Michael Gastegger, Philipp Marquetand

Molecular Dynamics with Neural Network Potentials

Chapter in book | Accepted manuscript (Postprint)

This version is available at <https://doi.org/10.14279/depositonce-10326>

Gastegger M., Marquetand P. (2020) Molecular Dynamics with Neural Network Potentials. In: Schütt K., Chmiela S., von Lilienfeld O., Tkatchenko A., Tsuda K., Müller KR. (eds) Machine Learning Meets Quantum Physics. Lecture Notes in Physics, vol 968. Springer, Cham.
https://doi.org/10.1007/978-3-030-40245-7_12

Terms of Use

Copyright applies. A non-exclusive, non-transferable and limited right to use is granted. This document is intended solely for personal, non-commercial use.

WISSEN IM ZENTRUM
UNIVERSITÄTSBIBLIOTHEK

Technische
Universität
Berlin

Molecular Dynamics with Neural-Network Potentials

Michael Gastegger and Philipp Marquetand

Abstract Molecular dynamics simulations are an important tool for describing the evolution of a chemical system with time. However, these simulations are inherently held back either by the prohibitive cost of accurate electronic structure theory computations or the limited accuracy of classical empirical force fields. Machine learning techniques can help to overcome these limitations by providing access to potential energies, forces and other molecular properties modeled directly after an accurate electronic structure reference at only a fraction of the original computational cost. The present text discusses several practical aspects of conducting machine learning driven molecular dynamics simulations. First, we study the efficient selection of reference data points on the basis of an active learning inspired adaptive sampling scheme. This is followed by the analysis of a machine-learning based model for simulating molecular dipole moments in the framework of predicting infrared spectra via molecular dynamics simulations. Finally, we show that machine learning models can offer valuable aid in understanding chemical systems beyond a simple prediction of quantities.

1 Introduction

Chemistry is – in large part – concerned with the changes that matter undergoes. As such, chemistry is inherently time-dependent and if we want to model such chemical processes, then a time-dependent approach is most intuitive. The corresponding techniques can be summarized as dynamics simulations. In particular molecular

Michael Gastegger

Technical University of Berlin, Machine Learning Group, Marchstr. 23, 10587 Berlin, Germany.
e-mail: michael.gastegger@tu-berlin.de

Philipp Marquetand

University of Vienna, Faculty of Chemistry, Institute of Theoretical Chemistry, Währinger Str. 17,
1090 Vienna, Austria. e-mail: philipp.marquetand@univie.ac.at

dynamics (MD) simulations – defined usually as treating the nuclear motion with Newton’s classical mechanics – are commonly used to ”mimic what atoms do in real life” [23]. Such simulations have become indispensable not only in chemistry but also in adjacent fields like biology and material science [1, 12].

An important ingredient of MD simulations are molecular forces, which determine how the nuclei move. For the sake of accuracy, it is desirable to obtain these forces from quantum-mechanical electronic structure calculations. However, such *ab initio* calculations are expensive from a computational perspective and hence only feasible for relatively small systems or short timescales. For larger systems (e.g. proteins) molecular forces are instead modeled by classical force fields, which are composed of analytic functions based on physical findings. As a consequence of these approximations, classical force fields are extremely efficient to compute but fail to reach the accuracy of electronic structure methods.

With the aim to obtain both a high accuracy and a fast evaluation, machine learning (ML) is employed to predict forces and an increasing number of research efforts are devoted to this idea. Since the forces are commonly evaluated as the derivative of the potential energy with respect to the coordinates, the learning of energies and forces are tightly connected and we often find the term ”machine learning potential” in the literature.

Possibly the first work to use ML of potentials combined with dynamics simulations appeared in 1995 by Blank *et al.* [8], where diffusion of CO on a Ni surface and H₂ on Si was modelled. A comprehensive overview over the earlier work in this field, where mostly neural networks were employed, can be found in the reviews [22, 4]. Later, also other ML methods like Gaussian approximation potentials [2] were utilized, diversifying the research landscape, which is reflected in various, more topical reviews, see e.g. [20, 28, 9, 6]. Today, the field has become so active, that a review would be outdated by tomorrow. Here instead, we give an example of what can be achieved when combining ML and MD.

In this chapter, we describe simulations of one of the most fundamental experiments to detect moving atoms, namely infrared spectra. The simulations utilize MD based on potentials generated with high-dimensional neural networks, a special ML architecture. We show in particular, how the training data can efficiently be gathered by an adaptive sampling scheme. Several practical aspects, tricks and pitfalls are presented. Special emphasis is put also on the prediction of dipole moments and atomic charges, which are necessary ingredients besides potential energies and forces for the calculation of infrared spectra from MD.

2 Methods

2.1 High-dimensional neural network potentials

High-dimensional neural network potentials (NNPs) are a type of atomistic ML potentials [7]. Atomistic potentials model the properties of a system based on the contributions of individual atoms due to their local chemical environment (Figure 1). In

Fig. 1 In a high-dimensional neural network potential, the local chemical environment of each atom is first encoded in a structural descriptor. Based on these descriptors, neural networks predict atomistic energy contributions, where different networks are used for the different chemical elements. Finally, the atomic energies are summed in order to recover the total energy of the system.

high-dimensional NNPs, atomic environments are represented via so-called atom-centered symmetry functions (ACSFs) [3]. Typically, ACSFs are radial and angular distribution functions, which account for rotational and translation invariances of the system. A radial symmetry function, for example, is a superposition of Gaussian densities:

$$G_i^{\text{rad}} = \sum_{j \neq i}^N e^{-\eta(r_{ij}-r_0)^2} f_{\text{cut}}(r_{ij}). \quad (1)$$

The sum includes all atoms j in vicinity of the central atom i . r_{ij} is the distance between i and j . η and r_0 are parameters which modulate the width and center of the Gaussian. A cutoff function f_{cut} ensures, that only the local chemical environment contributes to the ACSF. For a more in-depth discussion on ACSFs and their features, we refer to References [3, 5, 6, 16].

Based on the ACSF representation of each atom, an atomistic neural network then predicts the contribution of this atom to the global molecular property. Finally, these contributions are recombined via an atomistic aggregation layer in order to recover the target property. For NNPs which model the potential energy E of a system, this layer is usually chosen as a sum over the individual atomic energies E_i :

$$\tilde{E}^m = \sum_i^{N_m} \tilde{E}_i^m, \quad (2)$$

where N is the number of atoms in the molecule m . However, different aggregation layers can be formulated in order to model various properties, as will be discussed in the next section.

Due to the form of high-dimensional NNPs, expressions for analytic Cartesian derivatives of the model are readily available. Hence, NNPs provide access to energy conserving forces

$$\tilde{\mathbf{F}}_\alpha^m = - \sum_i^{N_m} \sum_d^{D_i^m} \frac{\partial \tilde{E}_i^m}{\partial G_d^m} \frac{\partial G_d^m}{\partial \mathbf{R}_\alpha^m}, \quad (3)$$

where $\tilde{\mathbf{F}}_\alpha^m$ and $\mathbf{R}_\alpha^m$ are the forces acting on atom α and its position, while D_i^m is the number of ACSFs centered on atom i . NNP forces can also be included into the training process by minimizing a loss function of the form

$$\mathcal{L} = \frac{1}{M} \sum_m^M (\tilde{E}^m - E^m)^2 + \frac{\vartheta}{M} \sum_m^M \frac{1}{3N_m} \sum_i^{N_m} \|\tilde{\mathbf{F}}_i^m - \mathbf{F}_i^m\|^2. \quad (4)$$

Here, M is the number of molecules present in the data set. ϑ controls the trade-off between fitting energies and forces, while $\mathbf{F}_i^m$ is the vector of Cartesian force components acting on atom i of molecule m . By using forces during training, $3N$ additional pieces of information are available for each molecule beside the potential energy. As a consequence, the overall number of reference computations required to construct an accurate NNP can be reduced significantly [26, 15].

2.2 Dipole model

In addition to energies and forces, atomistic aggregation layers can also be formulated to model various other molecular properties. One example is the dipole moment model introduced in Reference [13]. Here, the dipole moment $\boldsymbol{\mu}$ of a molecule is expressed as a system of atomic point charges, according to the relation

$$\tilde{\boldsymbol{\mu}} = \sum_i^N \tilde{q}_i \mathbf{r}_i. \quad (5)$$

$\tilde{q}_i$ is the charge located at atom i and $\mathbf{r}_i$ is the position vector of the atom relative to the molecules center of mass. The charges $\tilde{q}$ are modeled via atomistic networks and depend on the local chemical environment. However, these point charges are never learned directly, but instead represent latent variables inferred by the NNP dipole model during training, where the following loss function is minimized:

$$\mathcal{L} = \frac{1}{M} \sum_m^M (\tilde{Q}_m - Q_m)^2 + \frac{1}{3M} \sum_m^M \|\tilde{\boldsymbol{\mu}}_m - \boldsymbol{\mu}_m\|^2 + \dots \quad (6)$$

$\tilde{\boldsymbol{\mu}}$ is the expression for the dipole moment given in Equation 5 and $\boldsymbol{\mu}$ is the electronic structure reference, calculated as the expectation value of the dipole moment operator [19]. Note that the charges are not directly accessible from solving the Schrödinger equation but are usually obtained *a posteriori* from different *ad hoc* partitioning schemes [11]. The first term in the loss function introduces the additional constraint, that the sum of latent charges $\tilde{Q} = \sum_i \tilde{q}_i$ should reproduce the total charge Q of the molecule. Formulated in this way, the machine learning model depends only on quantum mechanical observables in the form of a molecule’s electrostatic moments (total charge and dipole moment), which are directly accessible by all electronic structure methods. Although the above formulation does not guarantee the conservation of total charge, it reduces the overall charge fluctuations to a minimum. The remaining deviations can then be corrected using simple rescaling or shifting schemes without loss of generality (see e.g. References [24] and [30]).

Expression 6 can easily be extended to include higher moments such as the quadrupole moment $\boldsymbol{\Theta}$, as was suggested in Reference [13]. In the context of the above model, $\boldsymbol{\Theta}$ takes the form

$$\tilde{\boldsymbol{\Theta}} = \sum_i \tilde{q}_i (3\mathbf{r}_i \otimes \mathbf{r}_i - \mathbf{1} \|\mathbf{r}_i\|), \quad (7)$$

where, $\mathbf{r}_i \otimes \mathbf{r}_i$ is the outer product of the Cartesian position vectors of atom i . However, it was found that the introduction of quadrupole moments offers no additional advantage, at least when modeling dipoles. Moreover, using an atomistic model for $\boldsymbol{\Theta}$ can be problematic for small molecules such as water, since the atom centered point charges are not able to resolve features of the charge distribution arising from e.g. the lone pair electrons of the oxygen.

It should be emphasized at this point, that the atomistic aggregation layers presented here are not restricted to a single type of machine learning architecture. They can be coupled with any model in a modular fashion, as long as it provides access to atomic contributions. This was for example done recently with the SchNet architecture in order to model dipole moment magnitudes [30].

2.3 Adaptive Sampling Scheme

Before a NNP can be used for simulations, its free parameters need to be determined by training on a suitable set of reference data. Typically, a set of reference molecules is chosen in a two-step process. First, the PES is sampled to obtain a representative set of molecular configurations. Afterwards, the quantum chemical properties of these structures (e.g. energies, forces, ...) are computed with an appropriate electronic structure method. The sampling can be performed in different ways, with molecular dynamics and normal mode sampling being only a few ex-

amples [5, 33]. However, a feature shared by most sampling methods is, that they either use approximate methods such as molecular force fields to guide the sampling or they perform all simulations at the final level of theory. Both approaches have drawbacks. In the first case, the PES regions explored with the lower level of theory need not correspond to regions relevant for the high level model (e.g. different molecular geometries). In the second case, the unfortunate scaling of accurate electronic structure methods limits either the regions of the PES that can be covered or the overall accuracy of the reference computations.

A solution to these issues is to use an adaptive sampling scheme, where the ML model itself selects new data points to improve its accuracy [13]. This approach is inspired by active learning techniques and proceeds as follows (Figure 2): First, a

Fig. 2 The adaptive sampling scheme starts by training a preliminary ensemble of NNPs on a small set of reference computations. This ensemble is then used to sample new configurations via e.g. molecular dynamics. For every sampled configuration, an uncertainty measure is computed. If this measure exceeds a predefined threshold, the simulation is stopped and a reference calculation is performed. The new data point is then added to the reference data and the next generation of NNPs is trained. This procedure is repeated in an iterative manner until a convergence threshold is reached.

crude NNP model is used to explore an initial region of the PES. During simulation, an uncertainty measure for the NNP predictions is monitored. If this measure exceeds a threshold, the sampling is stopped and electronic structure computations are performed for the corresponding configuration. The resulting data is then added to the reference set and used to update the NNP. Sampling is continued with the improved model. This procedure is carried out in an iterative fashion until the desired level of accuracy is reached.

One advantage of this scheme is that the NNP model used to guide the sampling closely resembles the electronic structure reference method for large stretches of coordinate space. Hence, similar regions of the PES will be explored (e.g. bond lengths) as if the simulations were carried out with the reference method exclusively.

In addition, by using the model uncertainty to determine when additional reference computations should be performed, only a small number of expensive calculations are necessary. Due to the simplicity of this scheme, it can easily be combined with different sampling methods, such as those in molecular dynamics, metadynamics or Monte-Carlo based approaches [17, 34].

Perhaps the most important ingredient for the above scheme is an appropriate uncertainty measure. Here, it is possible to make use of a trait of NNs in general and NNPs in particular. Two NNPs trained on the same reference data will agree closely for regions of the PES described well by both models. However, in regions sampled insufficiently the predictions of both models will diverge quickly. Using the disagreement between different models to guide data selection is a popular approach in active learning called query by committee [31]. Based on the above behavior, one can formulate the following uncertainty measure for NNPs:

$$\sigma_E = \sqrt{\frac{1}{\mathfrak{N}-1} \sum_{\mathfrak{n}} (\tilde{E}_{\mathfrak{n}} - \bar{E})^2}. \quad (8)$$

$\tilde{E}_{\mathfrak{n}}$ is the energy predicted by one of $\mathfrak{N}$ different NNPs and $\bar{E}$ is the average of all predictions. Hence, σ_E is the standard deviation of the different model predictions. Using an uncertainty measure of this form also has the following advantage: Since different NNPs are used to compute σ_E , they can be combined into an ensemble, where the prediction averages (e.g. $\bar{E}$) are used to guide PES exploration. The consequence is an improvement in the overall predictive accuracy of the ML approach at virtually no extra cost, due to error cancellation in the individual models.

3 Generation of Reference Data Sets

The following section discusses different practical aspects of the adaptive sampling scheme introduced above. After an investigation on the accuracy advantage offered by NNP ensembles, we study how frequently new reference computations are requested during a sampling run. Afterwards, the utility of different predicted properties as uncertainty measures for NNPs is analyzed. Finally, we introduce an extension to the standard sampling scheme, which improves overall sampling efficiency.

3.1 Accuracy of NNP Ensembles

In order to investigate the accuracy offered by ensembles of NNPs, we compare the predictions of ensembles containing up to five members to their respective electronic structure reference. This analysis is based on the protonated alanine tripeptide data set obtained in Reference [13], which also serves as a basis for several other studies in this text. The data set contains 718 different peptide configurations at the BLYP/def2-DZVP level of theory sampled with the scheme described above. Figure 3 shows the energy and force mean absolute errors (MAEs) and root mean squared errors (RMSEs) for the different ensembles. Even the combination of only two different models already leads to a marked decrease in the prediction error. Since ensembles thrive on a cancellation of random error fluctuations, this gain in

Fig. 3 Accuracy of ensemble predictions for molecular energies (blue) and forces (red) depending on the number of members. The computed error measures, MAE and RMSE, appear to decrease according to an $\frac{1}{\sqrt{\mathfrak{N}}}$ relation, where $\mathfrak{N}$ is the number of models in the ensemble. In all cases, the gain in accuracy is most pronounced when going from a single network to an ensemble of two.

accuracy is particularly pronounced for the RMSEs. An interesting observation is that the forces profit to a greater extent than the energies, with a reduction in the error by approximately $0.3 \text{ kcal mol}^{-1} \text{ Å}^{-1}$. This effect is expected to be of importance in the early stages of an adaptive sampling run, as the improved reliability of the model increases the likelihood that physically relevant configurations are sampled.

3.2 Choice of Uncertainty Measures

An important feature of atomistic NNPs is their ability to operate as fragmentation approaches, where they predict the energies of large molecules after being trained on only small fragments [14]. Hence, expensive reference computations never have to be performed for the whole system, but only for parts of it. This feature can be combined with the adaptive sampling scheme, as was for example done in Reference [13]. In this setup, the uncertainty is not measured for the whole molecule, but instead for atom-centered fragments. Reference computations are only performed for those fragments where the uncertainty exceeds a predefined threshold, thus reducing the computational cost of constructing an accurate NNP even further. How-

ever, the deviation of ensemble energies (Equation 8) can now no longer serve as the uncertainty measure.

Although substituting the total energies in Equation 8 by their atom-wise counterparts $\tilde{E}_i$ would in theory be a straightforward choice for an atomistic uncertainty estimate, it is not feasible in practice. Due to the way NNPs are constructed (see Equation 2), the partitioning of the total energy into latent atomic contributions is not unique. Hence, even if two NNPs yield almost identical predictions for the molecular energies, the distributions of atom-wise contributions can still differ significantly, as is shown for the alanine tripeptide in Figure 4. If e.g. the atomic ener-

Fig. 4 Distribution of atomic energies, forces and total energies as predicted for the alanine tripeptide by two NNP models (shown in red and blue). Although the NNP predictions agree well in the case of the total energies and atomic forces, the energy contributions of individual atoms vary dramatically between the models.

gies of carbon atoms are used to compute the uncertainty, large deviations will be encountered for all regions of the PES, no matter how well the global predictions agree. As a consequence, reference computations will be performed for a large fraction of encountered configurations, thus effectively negating the advantage offered by the adaptive sampling scheme.

The better alternative is to reformulate the above measure to instead use the forces acting on each atom:

$$\sigma_F^{(\alpha)} = \sqrt{\frac{1}{\mathfrak{N}-1} \sum_{\mathfrak{n}} \left\| \tilde{\mathbf{F}}_{\mathfrak{n}}^{(\alpha)} - \bar{\mathbf{F}}^{(\alpha)} \right\|^2}, \quad (9)$$

where $\tilde{\mathbf{F}}_n^{(\alpha)}$ is the force acting on atom α as predicted by model n of the ensemble. $\bar{\mathbf{F}}^{(\alpha)}$ is the average over all predictions. The measure $\sigma_F^{(\alpha)}$ has several advantages. Since it depends on the molecular forces it is purely atomistic. Moreover, due to how the forces are computed in NNPs (Equation 3), they are insensitive to the learned partitioning in a similar manner as the total energy. This property can be observed in Figure 4, where the distributions of forces acting on e.g. hydrogen and carbon atoms show a similar agreement between models as do the molecular energies, but not the atomic energies.

3.3 Frequency of Reference Computations

An important aspect of the adaptive sampling scheme is how frequently new electronic structure computations need to be performed. Figure 5 depicts the number of configurations added to the reference data set versus the total number of sampling steps. The studied molecule is an n-alkane ($\text{C}_{69}\text{H}_{140}$, see Figure inset) and the

Fig. 5 Number of configurations accumulated during an adaptive sampling run for the $\text{C}_{69}\text{H}_{140}$ n-alkane plotted against the number of molecular dynamics steps. New samples are added frequently during the early stages of the sampling, while almost no configurations are collected during the later stages.

sampling statistics were taken from the supporting information of Reference [13], obtained with a fragmentation procedure and the aforementioned force based uncertainty. As can be seen in the figure, there is a marked decrease in the number of electronic structure queries as the sampling progresses. Initially, new samples are added frequently, as the model explores the configuration space. More than half of the new samples are added within the first 2000 exploration steps. After this phase of determining a reliable first approximation of the electronic structure method, the

sampling process is dominated by fine-tuning the NNP ensemble. Now only samples corresponding to insufficiently described regions of the PES are collected, reducing the requirement for expensive reference computations significantly. The efficiency of this simple approach is remarkable insofar, as only 534 configurations are needed to obtain an accurate model of the n-alkane sporting 621 degrees of freedom. The final model achieves RMSEs of $0.09 \text{ kcal mol}^{-1}$ and $1.48 \text{ kcal mol}^{-1} \text{ \AA}^{-1}$ compared to the reference energies and forces of the fragments.

3.4 Adaptive Sampling with Multiple Replicas

A potential problem of the adaptive sampling scheme is its serial nature. Currently, only one point of data is collected after each sampling period. Since the NNPs need to be retrained every time the reference data set is extended, the resulting procedure can become time consuming in its later stages, especially for large and flexible molecules (e.g. the tripeptide in Reference [13]).

This problem can be overcome by introducing a parallel version of the adaptive sampling scheme, as outlined in Figure 6. Instead of simulating only a single system at a time,

Fig. 6 Parallel version of the adaptive sampling scheme. Individual adaptive sampling runs are carried out for different replicas of the system (e.g. different configurations). For each replica, configurations with high uncertainty are identified. Once samples have been collected for all replicas, reference computations are carried out and the NNP ensemble model is retrained. Afterwards, the replica simulations are continued with the new model. In this manner, the NNPs have to be retrained less frequently and different regions of the PES can be explored more efficiently.

multiple sampling runs are performed in parallel, each using a copy of the NNP ensemble trained in the previous cycle. These independent simulations can be replicas of the system under various conditions (e.g. different temperatures), a range of conformations or even different molecules. Sampling is once again carried out until divergence is reached for all parallel simulations. The high uncertainty configurations are then computed with the reference method and added to the training data. This setup reduces the frequency with which NNPs need to be retrained, while at the same time improving PES exploration. A potential drawback of this scheme is, that the collection of data points introduces periods of unproductivity, where some simulations are already finished while others are still running. However, this effect is negligible in praxis due to the high computational speed of the NNP models.

4 NNPs for Molecular Dynamics Simulations

Due to their combination of high accuracy and computational efficiency, NNPs are an excellent tool to accelerate MD simulations. A particularly interesting application is the computation of molecular spectra via the Fourier transform of different time autocorrelation functions [35]. Depending on the physical property underlying the autocorrelation function, different types of spectra can be obtained. One example are molecular infrared spectra, which can be modeled according to the relation:

$$I_{\text{IR}} \propto \int_{-\infty}^{+\infty} \langle \dot{\boldsymbol{\mu}}(\tau) \dot{\boldsymbol{\mu}}(\tau+t) \rangle_{\tau} e^{-i\omega t} dt, \quad (10)$$

where $\dot{\boldsymbol{\mu}}$ is the time derivative of the molecular dipole moment, τ is a time delay, ω is the vibrational frequency and t is the time.

The simulation of infrared spectra poses a particular challenge for machine learning techniques. Due to the dependence of Equation 10 on $\dot{\boldsymbol{\mu}}$, a reliable model of the molecular dipole moment $\boldsymbol{\mu}$ is needed in addition to the PES description provided by conventional NNPs. In the next sections, we will explore various aspects and the potential pitfalls associated with such models.

4.1 Machine Learning for Molecular Dipole Moments

A straightforward way to model dipole moments in the context of NNPs is to train individual atomic networks to reproduce quantum chemical partial charges. The molecular dipoles can then be obtained via the point charge model given in Expression 5, where the $\tilde{q}_i$ are modeled by environment-dependent networks. A similar strategy was e.g. used to model long-range electrostatic energies in Reference [24].

However, such a model suffers from the inherent inhomogeneity of the various charge partitioning schemes available for electronic structure methods. The pre-

dicted partial charges can differ dramatically between schemes and some of them fail at reproducing molecular dipole moments entirely [11]. Even when considering only those methods which yield partial charges consistent with the molecular dipole moment, a strong dependence on the type of partitioning can still be observed. Hirshfeld charges [18], for example, appear to work well in the setup described above, as was demonstrated in Reference [32]. Charges fit to the electrostatic potential (e.g. CHELPG [10]) on the other hand prove to be more problematic. To illustrate the issue at hand, Figure 7 shows the MD IR spectrum of single methanol molecule computed with a partial charge model based on the CHELPG method in comparison to the electronic structure reference. The partial charge spectrum shows several marked

Fig. 7 Infrared spectra of a methanol molecule in the gas phase computed via *ab initio* molecular dynamics (blue), as well as machine learned molecular dynamics using the dipole moment model introduced above (red) and a neural network model trained on CHELPG partial charges (gray). While the dipole moment model shows good agreement with the reference, the CHELPG model leads to erratic trends in peak positions and intensities.

differences from the reference. Small artificial peaks are introduced at 2100 and 3900 cm^{-1} respectively. Moreover, the intensity of several peaks (e.g. at 1400 cm^{-1} and 2800 cm^{-1}) is reproduced incorrectly. The most likely reason for these issues is the fitting procedure used to determine this particular type of reference charges. Since an independent least squares optimization is carried out for every molecular configuration, the obtained partial charges are not necessarily continuous with respect to incremental changes in the local environment of each atom. This makes it harder for the atomistic networks to learn a consistent charge assignment, leading to the erroneous behavior observed above.

A better approach is to incorporate the point charge model into the atomistic NNP architecture in the form of a dipole aggregation layer, as described in subsection 2.2 and Reference [13]. Instead of fitting to arbitrary partial charges, the model can now be trained directly on the molecular dipole moments, which are quantum

mechanical observables. In this manner, the need for choosing an appropriate partitioning scheme is eliminated. The inherent advantage of such a model can be seen in Figure 7, where it accurately reproduces the quantum chemical reference, although trained on the same set of configurations as the partial charge model.

4.2 Latent Partial Charges

A special feature of the above model is that it offers access to atomic partial charges. These charges are inferred by the NNP model based on the molecular electrostatic moments in a purely data driven fashion. Moreover, the charge models obtained with the above partitioning scheme depend on the local chemical environment of each atom. Hence, the charge distribution of the molecule can adapt to structural changes. Considering that partial charges are one of the most intuitive concepts in chemistry, the NNP latent charges represent an interesting analysis tool, e.g. for rationalizing reaction outcomes. In the following, we investigate how well the charges derived from the above dipole model agree with basic chemical insights.

One potential problem of atomistic properties (e.g. energies and charges) obtained via specialized aggregation layers is the extent with which the partitioning varies between different models. A good example are the atomic energies predicted by the tripeptide NNPs shown in Figure 4. Although the total energies agree well, the partitioning into atomic contributions is highly non-unique. Such a behavior is detrimental if the latent contributions should serve as an analysis tool. In order to investigate whether this phenomenon is also observed for the latent partial charges, a similar analysis is performed for two dipole moment models of the alanine tripeptide. The resulting partial charge distributions are compared in Figure 8. The latent

Fig. 8 Distribution of atomic partial charges predicted for the chemical elements present in the alanine tripeptide obtained with two dipole models (blue and red). Although differences between the two models are still present, the atomic charge distributions are better conserved than the atomic energies.

charges obtained with the dipole model are significantly better conserved than the atomic energies and only small deviations are found between different NNPs.

This trend appears to hold in general, as can be seen by repeating the above experiment for the QM9 database [27] containing approximately 130 000 small organic molecules. Figure 9 shows the distributions of partial charges and atomic energies of oxygen predicted for all molecules in QM9. In each case, five different dipole and

Fig. 9 Distribution of the atomic partial charges (blue) and energies (red) of oxygen predicted for all molecules in the QM9 database using different training set sizes (containing 5 000, 10 000, 50 000 and 100 000 datapoints). In each case, five dipole and energy models were trained on the magnitude of the dipole moment and the total electronic energy each. In contrast to the atomic energies, the partial charge distributions converge to similar values upon increasing the training set size.

energy models were trained on growing subsets of the database using an adapted version of ACSFs, so-called weighted ACSFs [16] composed of 22 radial and 10 angular symmetry functions. In case of the dipole moment models, we replaced the dipole vector $\vec{\mu}$ in equation 6 by its magnitude $|\vec{\mu}| = \|\sum_i^N \tilde{q}_i \mathbf{r}_i\|$, as only the latter is available in QM9. All models used atomistic networks with three layers of 100 neurons each and shifted softplus nonlinearities and were trained using the ADAM

algorithm [21] with a learning rate of 0.0001 (see References [29] and [30] for details). Compared to the atomic energies, the distributions of atomic partial charges are not only better conserved between models, but also show systematic convergence upon inclusion of additional data. The reason for this behavior is the geometry dependent term present in Equation 5, which introduces an additional constraint into the partitioning procedure. This term encodes information on the spatial distribution of molecular charge and is different for every molecule. Moreover, due to the statistical nature of the training procedure, the latent charge model has to be consistent across a wide range of molecules and configurations. The combination of both properties strongly limits the number of valid latent charge assignments. These results are encouraging and demonstrate that the NNP partial charges are indeed capable to capture aspects of the chemistry underlying a system. However, care should be taken when using the latent charges as a direct replacement of their quantum chemical counterparts, as the resulting partitionings – although well behaved – are still not uniquely determined. This can lead to undesirable effects when they are e.g. used to model long-range electrostatic interactions without further processing, as it can introduce inconsistencies into the predicted model energies and forces [36].

A final point of interest is the extent of fluctuations of the total charge observed during a typical molecular dynamics simulation (see section 2.2). Using the protonated alanine tripeptide as an example, the uncorrected total charge shows a standard deviation of 0.04 elementary charge units over a total of 150 ps of simulation time, while fluctuating around the expected average of 1.00. Hence, only minimal corrections are necessary to guarantee full charge conservation.

4.3 Electrostatic Potentials

Having ascertained the general reliability of the charge model, we now study how well the latent charge assignments agree with the predictions of electronic structure methods. In order to illustrate and compare different molecular charge distributions, we use partial charges to construct approximate electrostatic potentials (ESPs) of the form:

$$E(\mathbf{r}_0) = \sum_i^N \frac{q_i q_0}{\|\mathbf{r}_i - \mathbf{r}_0\|}, \quad (11)$$

where q_i and $\mathbf{r}_i$ are the partial charge and position vector of atom i . $\mathbf{r}_0$ is the position of a probe charge q_0 , which was set to $q_0 = +1$ in all experiments.

Figure 10 shows the pseudo ESPs obtained with latent and Hirshfeld partial charges. The latter have been chosen for their general reliability and widespread use. To assess, how well the latent predictions of the dipole model capture the charge redistribution associated with changes in the molecular geometry, two configurations of the protonated alanine tripeptide are modeled, with a hydrogen transferred from the N-terminal NH_3^+ group to the neighboring carbonyl moiety.

Fig. 10 Electrostatic potential surfaces of the alanine tripeptide based on Hirshfeld partial charges and latent partial charges yielded by the dipole model. The left-hand side shows a configuration protonated at the N-terminal NH_3^+ group, whereas the proton is situated on the adjacent carbonyl group in the right-hand side structure. Regions of negative charge are depicted in red, positively charged regions in blue.

In all cases, good agreement is found between the charge distributions predicted by the dipole moment model and the electronic structure equivalent. The latent charges are able to account for several important features, such as the localization of the positive charge of the molecule at the N-terminal NH_3^+ moiety in the first configuration, as well as its relocation towards the interior of the molecule after proton transfer. The model also accounts for the regions of negative charge expected for the carbonyl and carboxylic acid groups. These findings are remarkable insofar, as all this information on the electronic structure of the molecule is inferred purely from the global dipole moments, demonstrating the power of the partitioning scheme.

4.4 Geometry Dependence of Latent Charges

A final analysis is dedicated to the behavior of the latent dipole model charges under changes in the local chemical environment. As an example, we study the evolution of the partial charge of the proton during the proton transfer event occurring in the alanine tripeptide. Figure 11 shows the NNP partial charge attributed to the proton plotted against the reaction coordinate. The curves for Hirshfeld, Mulliken [25] and CHELPG charges are included for comparison. Several interesting effects can be observed.

First, the dipole model curve exhibits a minimum close to the transition state of the transfer reaction. Since the latent charges can be seen as a proxy of the local electron density, this result can be interpreted as follows: At the initial and final stages

Fig. 11 Changes in the partial charge of the proton during different stages of the proton transfer event. Shown are charges computed via different conventional charge partitioning schemes (Hirshfeld, CHELPG and Mulliken), as well as the latent charges predicted by the ML model.

of the transfer, the positive charge is located mainly at the proton itself. However, during the transfer and especially close to the transition regions, electron density is shared between the three participating atoms (O, N and proton). Hence, the positive charge is reduced for these configurations. This finding serves as an additional demonstration for the efficacy of the latent charge model. Although originally only conceived to model dipole moments, it is able to provide insights directly related to the electronic structure of the molecule at an atomistic resolution.

Second, Figure 11 illustrates the inherently different behavior found for various charge partitioning schemes. The Hirshfeld charges show a qualitatively similar curve to the machine learned charge model and are well behaved in general, supporting the results reported in Reference [32]. Mulliken and CHELPG charges on the other hand show completely different trends. The former are generally known for their unreliability, e.g. showing a strong dependence on the used basis set and large deviations when attempting to recover molecular dipole moments, hence the result is little surprising [11]. The counterintuitive behavior of the CHELPG charges serves as an additional confirmation for the effects observed in the methanol spectrum shown above (Figure 7). Given this general discrepancy between various partitioning schemes, the charges derived via the dipole moment model constitute a viable alternative: They reproduce molecular dipole moments accurately, are derived directly from quantum mechanical observables and capture the influence of structural changes on the molecular charge distribution.

5 Conclusion

We have presented how molecular dynamics (MD) simulations can benefit from machine learning (ML) potentials and provided some background for the imple-

mentation of this ML-MD approach. The first challenge during such a task is to efficiently gather enough training data in order to create a converged potential. An adaptive sampling scheme can serve for this purpose and the efficiency can be improved when using a) an ensemble of neural networks, b) an adequate uncertainty measure as selection criterion, and c) multiple replicas to parallelize the sampling.

As an ultimate test, experimental observables need to be calculated and compared to actual experimental results. In our case, infrared spectra are simulated, for which the neural networks do not only need to learn potentials and forces but also dipole moments and their atomistic counterparts, the atomic partial charges. If the latter are plotted in a geometry-dependent manner, e.g., along a reaction coordinate, these machine-learned charges provide insights directly related to the electronic structure of the molecule at an atomistic resolution. In this sense, machine learning can not only deliver potentials with supreme accuracy at compelling speed but also offer valuable insights beyond a simple prediction of quantities.

Despite this positive picture, many challenges remain, e.g., generalizing the machine learning models – ideally for all possible substances, different kinds of molecules and materials alike – or extending the range of properties to be learned. Possibly the biggest challenge is however to find a universally valid electronic structure method necessary for the generation of high-fidelity training data.

Acknowledgements M.G. was provided financial support by the European Unions Horizon 2020 research and innovation program under the Marie Skłodowska-Curie grant agreement NO 792572. The computational results presented have been achieved in part using the Vienna Scientific Cluster (VSC). We thank J. Behler for providing the RuNNer code.

References

1. Allen, M. P. and Tildesley, D. J. (1987) *Computer Simulation of Liquids*. Oxford University Press, Oxford.
2. Bartók, A. P., Payne, M. C., Kondor, R., and Csányi, G. (2010) Gaussian approximation potentials: The accuracy of quantum mechanics, without the electrons. *Phys. Rev. Lett.*, **104**, 136403.
3. Behler, J. (2011) Atom-centered symmetry functions for constructing high-dimensional neural network potentials. *J. Chem. Phys.*, **134**, 074106.
4. Behler, J. (2011) Neural network potential-energy surfaces in chemistry: a tool for large-scale simulations. *Phys. Chem. Chem. Phys.*, **13**, 17930–17955.
5. Behler, J. (2015) Constructing high-dimensional neural network potentials: A tutorial review. *Int. J. Quantum Chem.*, **115**, 1032–1050.
6. Behler, J. (2017) First principles neural network potentials for reactive simulations of large molecular and condensed systems. *Angew. Chem. Int. Edit.*, **56**, 12828–12840.
7. Behler, J. and Parrinello, M. (2007) Generalized Neural-Network Representation of High-Dimensional Potential-Energy Surfaces. *Phys. Rev. Lett.*, **98**, 146401.
8. Blank, T. B., Brown, S. D., Calhoun, A. W., and Doren, D. J. (1995) Neural network models of potential energy surfaces. *J. Chem. Phys.*, **103**, 4129–4137.
9. Botu, V., Batra, R., Chapman, J., and Ramprasad, R. (2017) Machine learning force fields: Construction, validation, and outlook. *J. Phys. Chem. C*, **121**, 511–522.

10. Breneman, C. M. and Wiberg, K. B. (1990) Determining atom-centered monopoles from molecular electrostatic potentials. the need for high sampling density in formamide conformational analysis. *J. Comput. Chem.*, **11**, 361–373.
11. Cramer, C. J. (2004) *Essentials of Computational Chemistry*. Wiley, 2nd edn.
12. Frenkel, D. and Smit, B. (2001) *Understanding Molecular Simulation*. Academic Press.
13. Gastegger, M., Behler, J., and Marquetand, P. (2017) Machine learning molecular dynamics for the simulation of infrared spectra. *Chem. Sci.*, **8**, 6924–6935.
14. Gastegger, M., Kauffmann, C., Behler, J., and Marquetand, P. (2016) Comparing the accuracy of high-dimensional neural network potentials and the systematic molecular fragmentation method: A benchmark study for all-trans alkanes. *J. Chem. Phys.*, **144**, 194110.
15. Gastegger, M. and Marquetand, P. (2015) High-dimensional neural network potentials for organic reactions and an improved training algorithm. *J. Chem. Theory Comput.*, **11**, 2187–2198.
16. Gastegger, M., Schwiedrzik, L., Bittermann, M., Berzsenyi, F., and Marquetand, P. (2018) wACSF – Weighted atom-centered symmetry functions as descriptors in machine learning potentials. *J. Chem. Phys.*, **148**, 241709.
17. Herr, J. E., Yao, K., McIntyre, R., Toth, D. W., and Parkhill, J. (2018) Metadynamics for training neural network model chemistries: A competitive assessment. *J. Chem. Phys.*, **148**, 241710.
18. Hirshfeld, F. (1977) Bonded-atom fragments for describing molecular charge densities. *Theor. Chim. Acta*, **44**, 129–138.
19. Jensen, F. (2007) *Introduction to Computational Chemistry*. John Wiley & Sons, Ltd, 2 edn.
20. Jiang, B., Li, J., and Guo, H. (2016) Potential energy surfaces from high fidelity fitting of ab initio points: the permutation invariant polynomial - neural network approach. *Int. Rev. Phys. Chem.*, **35**, 479–506.
21. Kingma, D. P. and Ba, J. (2014) Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*.
22. Latino, D. A. R. S., Fartaria, R. P. S., Freitas, F. F. M., Aires-De-Sousa, J., and Silva Fernandes, F. M. S. (2010) Approach to potential energy surfaces by neural networks. A review of recent work. *Int. J. Quantum Chem.*, **110**, 432–445.
23. Marx, D. and Hutter, J. (2009) *Ab Initio Molecular Dynamics: Basic Theory and Advanced Methods*. Cambridge University Press, Cambridge.
24. Morawietz, T., Sharma, V., and Behler, J. (2012) A neural network potential-energy surface for the water dimer based on environment-dependent atomic energies and charges. *J. Chem. Phys.*, **136**, 064103.
25. Mulliken, R. S. (1955) Electronic population analysis on LCAO-MO molecular wave functions. I. *J. Chem. Phys.*, **23**, 1833–1840.
26. Pukrittayakamee, A., Malshe, M., Hagan, M., Raff, L. M., Narulkar, R., Bukkapatnam, S., and Komanduri, R. (2009) Simultaneous fitting of a potential-energy surface and its corresponding force fields using feedforward neural networks. *J. Chem. Phys.*, **130**, 134101.
27. Ramakrishnan, R., Dral, P. O., Rupp, M., and Von Lilienfeld, O. A. (2014) Quantum chemistry structures and properties of 134 kilo molecules. *Scientific data*, **1**, 140022.
28. Ramakrishnan, R. and von Lilienfeld, O. A. (2017) *Machine Learning, Quantum Chemistry, and Chemical Space*, chap. 5, pp. 225–256. Rev. Comput. Chem., John Wiley & Sons, Inc.
29. Schütt, K., Kessel, P., Gastegger, M., Nicoli, K., Tkatchenko, A., and Müller, K.-R. (2018) Schnetpack: A deep learning toolbox for atomistic systems. *J. Chem. Theory Comput.*, **15**, 448–455.
30. Schütt, K. T., Gastegger, M., Tkatchenko, A., and Müller, K.-R. (2018) Quantum-chemical insights from interpretable atomistic neural networks. *arXiv:1806.10349 [physics.comp-ph]*.
31. Seung, H. S., Oppen, M., and Sompolinsky, H. (1992) Query by committee. *Proceedings of the fifth annual workshop on Computational learning theory*, pp. 287–294, ACM.
32. Sifain, A. E., Lubbers, N., Nebgen, B. T., Smith, J. S., Lokhov, A. Y., Isayev, O., Roitberg, A. E., Barros, K., and Tretiak, S. (2018) Discovering a transferable charge assignment model using machine learning. *J. Phys. Chem. Lett.*, **9**, 4495–4501.
33. Smith, J. S., Isayev, O., and Roitberg, A. E. (2017) ANI-1: an extensible neural network potential with DFT accuracy at force field computational cost. *Chem. Sci.*, **8**, 3192–3203.

34. Smith, J. S., Nebgen, B., Lubbers, N., Isayev, O., and Roitberg, A. E. (2018) Less is more: Sampling chemical space with active learning. *J. Chem. Phys.*, **148**, 241733.
35. Thomas, M., Brehm, M., Fligg, R., Vohringer, P., and Kirchner, B. (2013) Computing vibrational spectra from ab initio molecular dynamics. *Phys. Chem. Chem. Phys.*, **15**, 6608–6622.
36. Yao, K., Herr, J. E., Toth, D., Mckintyre, R., and Parkhill, J. (2018) The tensormol-0.1 model chemistry: a neural network augmented with long-range physics. *Chem. Sci.*, **9**, 2261–2269.