

TOPICAL REVIEW • OPEN ACCESS

Strategies for the construction of machine-learning potentials for accurate and efficient atomic-scale simulations

To cite this article: April M Miksch *et al* 2021 *Mach. Learn.: Sci. Technol.* **2** 031001

View the [article online](#) for updates and enhancements.

You may also like


- [Revisiting the stable structures of gold clusters: Au_n \(n = 16–25\) by artificial neural network potential](#)
Yibo Guo, Xue Wu and Jie Fu
- [An Improved Artificial Neural Network Design for Face Recognition utilizing Harmony Search Algorithm](#)
Maryam Mahmood Hussein, Ammar Hussein Mutlag and Hussain Shareef
- [ECoPANN: A Framework for Estimating Cosmological Parameters Using Artificial Neural Networks](#)
Guo-Jian Wang, Si-Yao Li and Jun-Qing Xia



TOPICAL REVIEW

OPEN ACCESS

Strategies for the construction of machine-learning potentials for accurate and efficient atomic-scale simulations

RECEIVED
22 January 2021REVISED
20 April 2021ACCEPTED FOR PUBLICATION
30 April 2021PUBLISHED
14 July 2021April M Miksch^{1,*} , Tobias Morawietz² , Johannes Kästner¹ , Alexander Urban^{3,4} 
and Nongnuch Artrith^{3,4,*} ¹ Institute for Theoretical Chemistry, University of Stuttgart, 70569 Stuttgart, Germany² Bayer AG, Pharmaceuticals, R&D, Digital Technologies, Computational Molecular Design, 42096 Wuppertal, Germany³ Department of Chemical Engineering, Columbia University, New York, NY 10027, United States of America⁴ Columbia Center for Computational Electrochemistry, Columbia University, New York, NY 10027, United States of America

* Authors to whom any correspondence should be addressed.

E-mail: miksch@theochem.uni-stuttgart.de and nartrith@atomistic.net**Keywords:** artificial neural networks, interatomic potentials, tutorial, active learningOriginal Content from this work may be used under the terms of the [Creative Commons Attribution 4.0 licence](https://creativecommons.org/licenses/by/4.0/).

Any further distribution of this work must maintain attribution to the author(s) and the title of the work, journal citation and DOI.

**Abstract**

Recent advances in machine-learning interatomic potentials have enabled the efficient modeling of complex atomistic systems with an accuracy that is comparable to that of conventional quantum-mechanics based methods. At the same time, the construction of new machine-learning potentials can seem a daunting task, as it involves data-science techniques that are not yet common in chemistry and materials science. Here, we provide a tutorial-style overview of strategies and best practices for the construction of artificial neural network (ANN) potentials. We illustrate the most important aspects of (a) data collection, (b) model selection, (c) training and validation, and (d) testing and refinement of ANN potentials on the basis of practical examples. Current research in the areas of active learning and delta learning are also discussed in the context of ANN potentials. This tutorial review aims at equipping computational chemists and materials scientists with the required background knowledge for ANN potential construction and application, with the intention to accelerate the adoption of the method, so that it can facilitate exciting research that would otherwise be challenging with conventional strategies.

1. Introduction

First-principles-based atomic scale simulations, for example using density-functional theory (DFT) [1–6], can predict many materials properties with quantitative accuracy [7–14]. However, they are usually limited to small atomic structures with less than 1000 atoms and time scales on the order of nanoseconds, owing to the high computational demand and polynomial scaling with the system size. During the last decade, a family of methods for accelerating first-principles sampling based on machine learning (ML) has emerged [15–21], which holds promise to overcome these limitations. ML regression models, usually based on artificial neural networks (ANNs) [22, 23] or Gaussian process regression (GPR) [24], are trained to interpolate the outcomes from first-principles calculations, so that the trained ML model can be used as computationally efficient drop-in replacements for the original method.

Predicting the preferred atomic structure at specific thermodynamic conditions and its evolution over a certain period of time requires a description of the relative energy of atomic arrangements, i.e. the potential energy surface (PES). First-principles methods, such as DFT or quantum-chemical methods based on wavefunction theory [25–30], approximate the PES based on the interactions of electrons and atomic nuclei arising from the laws of quantum mechanics.

Most ML potential (MLP) approaches do not consider electronic interactions explicitly but instead approximate the PES as a function of the atomic positions only. For many modeling applications, this general strategy allows MLPs to deliver the accuracy of the reference method at a computational cost that is orders of magnitude lower and scales only linearly with the number of atoms. Owing to the success of early

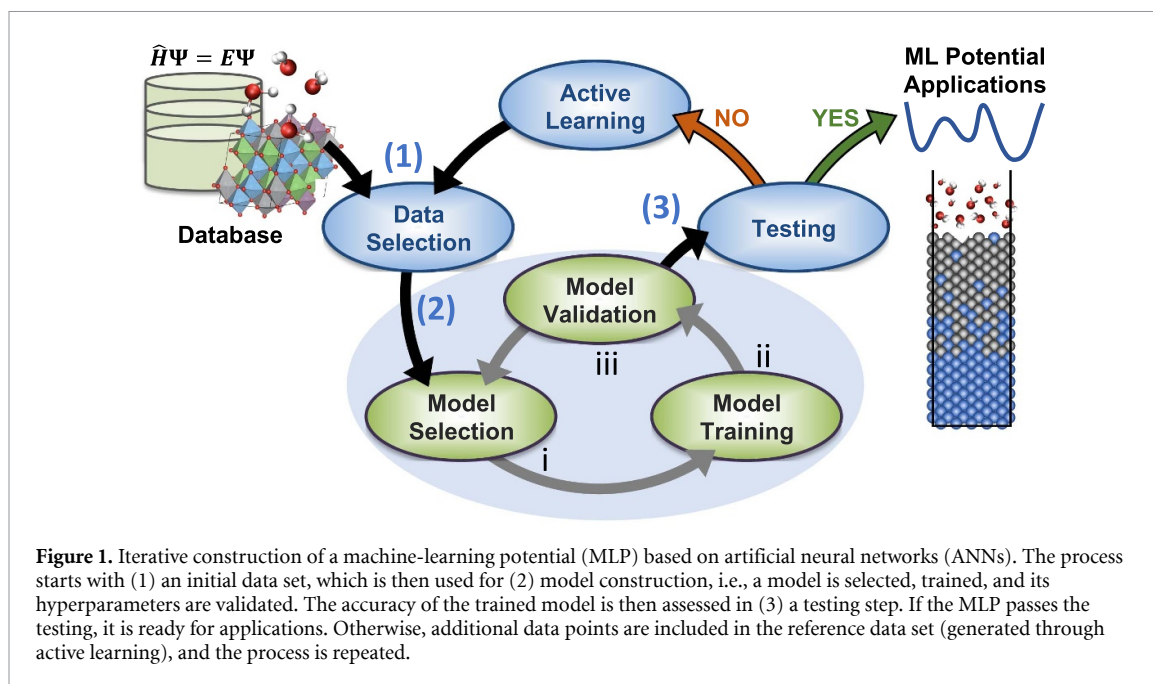


Figure 1. Iterative construction of a machine-learning potential (MLP) based on artificial neural networks (ANNs). The process starts with (1) an initial data set, which is then used for (2) model construction, i.e., a model is selected, trained, and its hyperparameters are validated. The accuracy of the trained model is then assessed in (3) a testing step. If the MLP passes the testing, it is ready for applications. Otherwise, additional data points are included in the reference data set (generated through active learning), and the process is repeated.

ANN-based Behler-Parrinello MLPs [23, 31] and GPR-based Gaussian approximation potentials (GAPs) [24, 32], the number of MLP methods proposed in the literature has been rapidly growing: examples include MLPs based on GPR and other kernel-based methods, such as kernel ridge regression [33–36], moment tensor potentials [37, 38], graph-networks using message passing [39–47], spectral neighbor analysis potentials (SNAPs) [48, 49], and other ANN-based approaches [50–53]. We emphasize that this list is not exhaustive and also does not include the various ML methods for atomistic modeling that cannot be considered interatomic potentials [54–70]. For a comparison of different MLP methods we refer the reader also to perspectives and reviews by Behler [15], Deringer [71], Mueller [17], Noé [18], and Unke [19, 21].

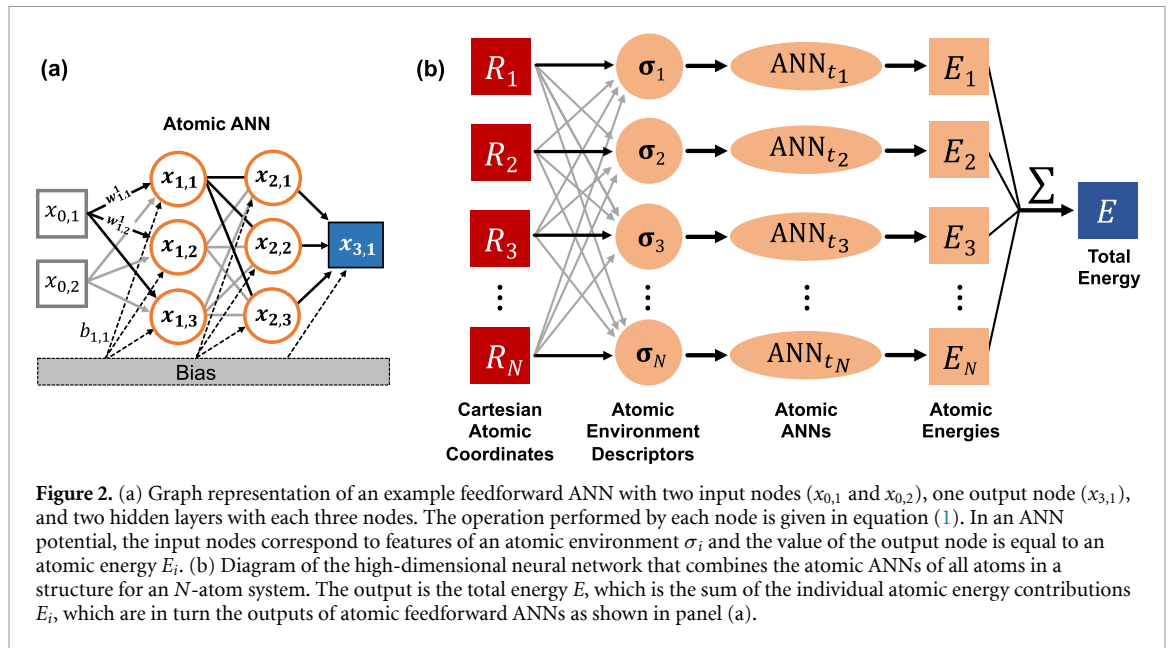
Although the ML regression models can be used in simulations in the same fashion as conventional interatomic potentials (*force fields*) [72, 73], the construction and applicability range of MLPs is significantly different. Here, we review and discuss the practical aspects of constructing and validating MLPs in the form of a tutorial with concrete examples. To be as specific as possible, the tutorial focuses on ANN-based MLPs (ANN potentials), although many aspects of the discussion on data selection (section 2) and active learning (section 5) apply to other MLP methods as well. The sections on model selection (section 3) and training/validation (section 4) are mostly specific to ANN potentials.

The construction of ANN potentials and other MLPs is centered around the compilation of reference data sets with atomic structures and their corresponding first principles energies, and potentially further information such as interatomic forces and atomic charges. Unlike many conventional potentials with functional forms derived from physical approximations, MLPs are usually not capable of extrapolating to atomic structures and compositions that lie outside of the PES region described by the reference data. The reference data set therefore needs to span the entire structural and chemical space that the MLP must represent for the intended application range, while it should include as few unnecessary or redundant data points as possible. In practice, MLP construction and the compilation of reference data is therefore (typically) an iterative process, shown in figure 1, that involves:

- (a) data collection,
- (b) model construction, and
- (c) testing,

which is repeated until the MLP passes the testing step with the desired accuracy. In each iteration, the reference data set is extended through an *active learning* scheme.

As also hinted at in figure 1, the model construction consists itself of three iterative steps: (i) *Model selection* is the process of deciding the type, *complexity*, and other *hyperparameters* of the ML model; (ii) *training* is the optimization of the adjustable model parameters to best reproduce the reference data set, as measured by a *loss function*; and (iii) in the *validation* step, *over-* or *underfitting* is detected, and, if necessary, the process is repeated from step (i) with adjusted model complexity and hyperparameters.



In this work, we discuss each of the steps outlined in figure 1 from a practical perspective. The next section first provides a brief introduction to the ANN potential method, which is followed by separate sections on data selection, the individual steps of model construction, testing, and active learning. In the final discussion section, current limitations and advanced techniques of the ANN potential method are considered.

1.1. High-dimensional ANN potentials

ANNs are a class of mathematical functions that can be represented by graphs which resemble networks of nodes (*artificial neurons*) that calculate the weighted sum of multiple input values and apply an *activation function* to the result. The operation performed by a single node can be expressed as

$$x_{i,j} = f_a^i \left(\sum_k^{N_{i-1}} w_{k,j}^i x_{i-1,k} + b_{i,j} \right) \quad (1)$$

where the output $x_{i,j}$ is the value of the j th node in the i th layer of the ANN, the input value $x_{i-1,k}$ is the k th node in the previous layer ($i-1$), $w_{k,j}^i$ is the *weight* of the input value, $b_{i,j}$ is an additional *bias weight* that is always added to the input irrespective of the input values, and f_a^i is an *activation* (or *transfer*) function. Equivalently, equation (1) can also be expressed using matrix-vector operations

$$\vec{x}_i = f_a^i \left(\mathbf{W}_i \vec{x}_{i-1} + \vec{b}_i \right), \quad (2)$$

where the vectors $\vec{x}_{i-1} = (x_{i-1,1}, \dots, x_{i-1,N_{i-1}})^T$ and $\vec{x}_i = (x_{i,1}, \dots, x_{i,N_i})^T$ contain the values of all nodes in the input and output *layers*, respectively, $(\mathbf{W}_i)_{k,j} = w_{k,j}^i$ is the weight matrix, and $(\vec{b}_i)_j = b_{i,j}$. ANN potentials are based on *feedforward* ANNs in which layers of nodes are connected such that values are passed layer-by-layer from an input layer through one or more hidden layers to an output layer. A graph representation of an example feedforward ANN is shown in figure 2(a). The number of layers, the number of nodes per layer, and the choice of activation function are hyperparameters that need to be chosen in the model selection step and are discussed in section 3.2. Note that the input dimension (number of nodes in the input layer) and output dimension (number of nodes in output layer), i.e. the dimensions of the domain and co-domain of the ANN function, are fixed for a given ANN. ANN training is the process of optimizing the weight parameters $\{w_{k,j}^i\}$ and $\{b_{i,j}\}$ such that the target values of the reference data set are reproduced as well as possible (see section 4). For a more thorough introduction to ANNs, we refer the reader to previous literature [50] and standard textbooks [74].

In principle, PESs can be directly represented by ANNs, feeding the atomic coordinates (in form of internal coordinates) into the input layer and producing the potential energy as the sole value of the output layer [75]. While useful for many applications, such direct ANN-PES models are limited to a fixed number of atoms and do not automatically reflect the physical invariances of the potential energy with respect to

rotation and translation of the entire atomic structure and the exchange of equivalent atoms. As such, direct ANN-PES models are not comparable to conventional interatomic potentials.

To leverage the flexibility of ANNs for the construction of actual, reusable interatomic potentials, Behler and Parrinello proposed an alternative approach [23] that was later extended to multiple chemical species by Artrith, Morawietz, and Behler [76], in which the total energy $E(\sigma)$ of an atomic structure σ is decomposed into atomic energy contributions E_i

$$E(\sigma) \approx \sum_i^{\text{atoms}} E_i(\sigma_i), \quad (3)$$

with $\sigma_i = \{\vec{R}_j, t_j \text{ for } |\vec{R}_j - \vec{R}_i| \leq R_{\text{cut}}\}$

where σ_i represents the *local atomic environment* of atom i that contains the positions \vec{R}_j and species t_j of all atoms within a cutoff distance R_{cut} from atom i , including the position \vec{R}_i and type t_i of atom i itself. In the high-dimensional ANN potential method by Behler and Parrinello, the atomic energies $E_i(\sigma_i)$ are predicted by species-specific ANNs. A graph representation of such a high-dimensional ANN potential is shown in figure 2(b).

As expressed in equation (3), the ANN potential does not yet incorporate the invariances of the potential energy, since the input of the atomic ANNs are still atomic (Cartesian) coordinates. Even worse, the number of atoms within the local atomic environment is generally structure dependent, but the input dimension of ANNs is fixed, which would render the atomic ANNs essentially not transferable to other structures. These limitations can be removed by representing the atomic coordinates within local atomic environments with a fixed number of *features* that have the same invariances as the potential energy. Once the atomic species are also encoded, a *fingerprint* $\tilde{\sigma}_i$ of the local atomic environment σ_i is obtained that can be used as input for the atomic ANNs, so that equation (3) can be written as

$$E(\sigma) \approx E_{\text{ANN}}(\sigma) = \sum_t^{\text{atom types}} \sum_i^{\text{atoms of type } t} \text{ANN}_t(\tilde{\sigma}_i), \quad (4)$$

where ANN_t is the atomic ANN for atoms of type t .

It is important to note that the high-dimensional ANN potential method as introduced by Behler and Parrinello uses ANNs to represent atomic energies, even though *the training targets are the total energies*. Atomic energies are not uniquely defined and are not a quantum-mechanical observable. ANN potential training, i.e. the weight optimization problem, can be expressed completely in terms of the total energies without knowledge of the *atomic energies*, as discussed further in section 4.1.3. It is possible to come up with rigorous (but non-unique) definitions of the atomic energies and use those as targets for training [77], but such schemes require additional processing of first principles calculations and are not considered in the following.

Further note that the atomic-energy approach of equation (4) can only work with good accuracy if the total energy is determined fully by short-ranged interactions. Long- and intermediate-ranged interactions, such as electrostatic (*Coulomb*) and dispersive (*London* or *van der Waals*) interactions need to be accounted for separately. Extensions of the high-dimensional ANN potential method to electrostatic and dispersive interactions have been developed and are briefly discussed in section 6.

2. Data selection

Since the mathematical form of ANN potentials and other types of MLPs is unconstrained and not derived from physical approximations, MLPs are poor at extrapolating to atomic structures or compositions that are very different from those included in the data that the MLP was trained on. The lack of extrapolation capabilities is, in fact, a general property of ANNs [78]. The quality of an ANN potential therefore depends strongly on the reference data set that it is trained on:

- an MLP's *accuracy* for the prediction of materials or molecular properties cannot exceed that of the reference method, and
- the *transferability* (the ability to *generalize*) of an MLP is determined by the structural and chemical space that is represented in the reference data set.

Incorrect data points, noise, and redundant data can further impede the MLP training process. Data selection is therefore of great importance for the construction of accurate and transferable MLPs [79, 80].

The atomic structures and compositions within the reference data set determine the *feature space*. The *target space* depends on the types of derived physical properties that are present in the data set and can be represented by the selected model (see section 3). In addition to total energies, other quantifiable physical properties, such as atomic charges [46, 76, 81], electronegativities [82–84], molecular dipole moments [63, 85], and atomic forces or higher-order derivatives [76, 86, 87] can in principle be included in the reference data set.

In this section, we discuss strategies for the compilation of *initial* reference data sets, i.e. data sets generated before any model construction or testing has occurred. The refinement of reference data in subsequent iterations of the process in figure 1 is discussed in the context of active learning approaches in section 5.

While data selection is crucial for any kind of MLP, there are some specifics that apply to ANN potentials: in their conventional form, ANN potentials can be most easily trained on energies and forces (see also section 4), but it is challenging to include other target properties of interest (such as higher order derivatives of the potential energy) as well. This is an area where other MLP methods could be better suited. The computational cost of evaluating ANN potentials does not increase with the size of the reference data set, so that ANN potentials can be trained on very large reference data sets containing millions of data points [51]. This is not generally the case for MLP methods, and depending on the method it can be beneficial to reduce the size of the reference data set.

2.1. Recipe: generation of initial reference data sets

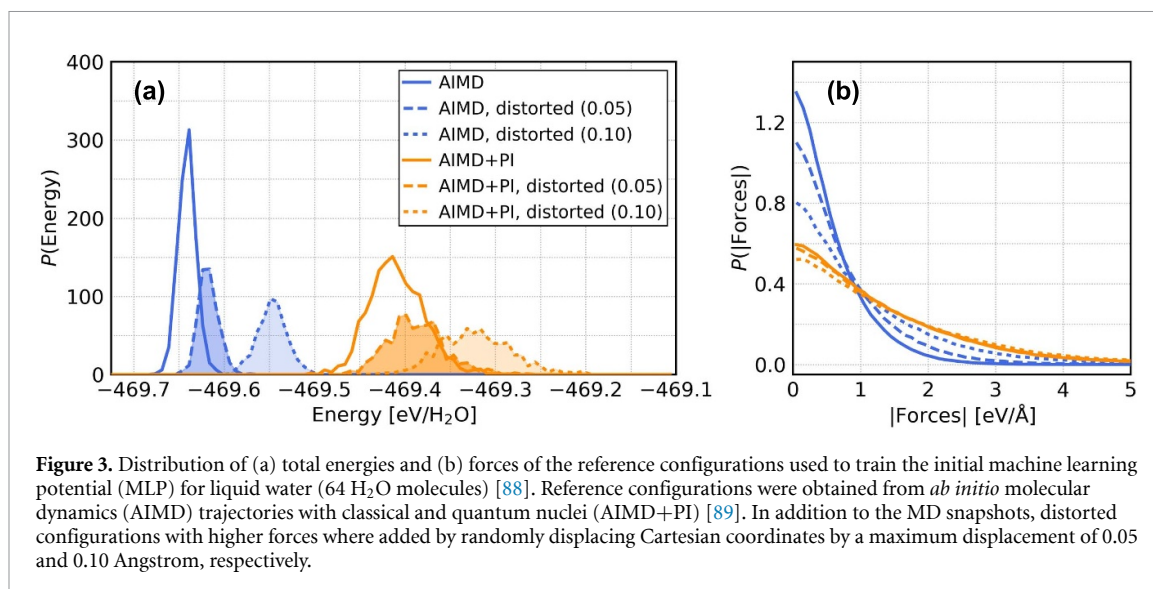
The initial reference data set is used to *kick off* the iterative refinement of the MLP shown in the flow chart of figure 1. The data set should already include structures that are close to those relevant for the eventual application of the potential. An initial data set could, for example, comprise of

- (a) relevant ideal atomic structures from databases, e.g. ideal crystal structures from crystallographic databases;
- (b) structures that were *derived* from the ideal structures by modifying the *atomic positions*, e.g. by displacing atoms;
- (c) derived structures that were generated by altering the *lattice parameters* or by scaling the ideal structures; and
- (d) derived structures that exhibit relevant defects, e.g. point defects, substitutional disorder, etc.

Including distorted structures is particularly important, so that the initial data set contains structures with bond length that are significantly shorter and longer than those in the ideal structures. While compiling the initial data set, it is important to keep an eye on its energy distribution: (a) structures that are so high in relative energy that they are realistically never encountered in the planned simulations should not be included in the data set, (b) the data set should not exhibit gaps in energy space, i.e. the energies between the most and least stable structures should be relatively evenly represented in the data set. Note that the maximal reasonable relative energy depends on the specific application that the potential is going to be used for.

2.2. Example: an initial reference data set for liquid water

MLP-based simulation have previously been shown to achieve high accuracy for both liquid and solid water [12, 80, 88, 90–92]. Here, we illustrate some of the considerations that go into the selection of an initial data set using the example of an MLP for liquid water which was applied to the calculation of vibrational spectra across the full liquid temperature range [88, 91]. Figures 3(a) and (b) show the distribution of interatomic energies and forces in an initial reference data set for the liquid water MLP. In this particular case, the data set was obtained by selecting periodic water structures (containing 64 H₂O molecules) from DFT-based *ab initio* molecular dynamics (AIMD) simulations at ambient conditions, in which the thermal motion of the atoms was used to sample the relevant structure space [89]. In addition to configurations from AIMD simulations with classical nuclei, the initial data set also contains snapshots from AIMD trajectories with *quantum nuclei* as described by the path integral (PI) method [93] which are located in a higher energy region of the PES (shown in orange in figure 3). Furthermore, distorted structures were generated by randomly displacing atoms from the original structures (obtained from equally spaced snapshots of the AIMD and AIMD+PI trajectories) with maximum displacements of 0.05 Å and 0.10 Å, respectively, to further increase the structural diversity in the initial reference data set. The initial data set comprising of a total of 5369 atomic structures was then used for the construction of an initial ANN potential. As seen in figure 3, the initial structures fully cover a large energy range of 500 meV/H₂O (the thermal energy $k_B T$, where k_B is Boltzmann's constant, at 300 K is around 26 meV per degree of freedom) which can be expected to be sufficient for performing stable MD simulations even at temperatures above ambient conditions with the ANN potential trained to this initial data set.



3. Model selection

Once an initial data set has been compiled, the next step in the flow chart of figure 1 is the selection of a model to represent the data. This means, in the case of ANN potentials, to decide on (A) the descriptor that is used for the featurization of local atomic environments, i.e. the *fingerprint* $\tilde{\sigma}_i$ in equation (4), and (B) the ANN hyperparameters, for example, the number of ANN layers, the number of nodes per layer, and the type of activation function.

3.1. Recipe: featurization of local atomic environments

As discussed in section 1.1, to be suitable as input for an ANN, the coordinates and atomic species of the atoms within the local atomic environment σ_i of an atom i need to be transformed into a feature vector with fixed length that should also be invariant with respect to rotation and translation of the atomic structure and the exchange of equivalent atoms. The amount of detail captured by this *descriptor* also determines the ability of the ANN potential to distinguish between different atomic structures: If the descriptor is too approximate, different atomic structures might yield the same feature vector. Hence, the choice of descriptor is crucial for the accuracy of the ANN potential.

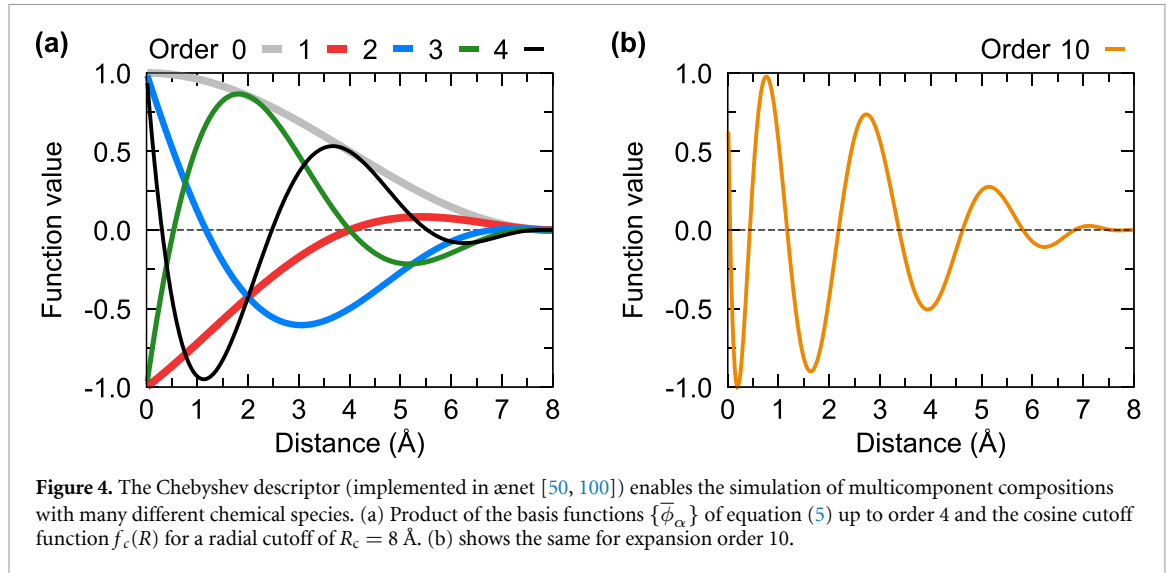
The choices to be made are (a) the type of descriptor that is used for the featurization of the local atomic structure and compositions, (b) the resolution of the descriptor, and (c) the radial cutoff, i.e. the size of the local atomic environment.

3.1.1. Selecting a descriptor

Various descriptors have been proposed in the literature [94], most of which are derived from basis-set expansions of either the atomic density of the local atomic environment [24, 32, 95–98], the radial and angular distribution functions (RDF and ADF) and higher-order distribution functions within the local atomic environment [23, 51, 99–103], or directly the local PES [37, 104]. In addition to differences in the features for the geometry of the local atomic environment, the above descriptors also differ in their approach for encoding chemical species. We note that the above list of descriptors is not exhaustive, and the development of new methods is currently an active field of research. Several of the expansion-based descriptors are available in open-source libraries [105, 106]. As an alternative to expansion-based descriptors for *hand-crafted* feature construction, an invariant representation of the atomic coordinates can also be machine learned [42, 45, 46, 107–109].

For the purpose of ANN potentials, an ideal descriptor

- exhibits the symmetries of the potential energy,
- requires minimal manual fine-tuning,
- provides a parameter for adjusting its resolution,
- is continuous and differentiable, so that analytical forces can be obtained,
- is computationally efficient, and
- does not scale with the number of chemical species.



Here, we discuss one specific choice of descriptor that fulfills the above criteria, has been frequently used for the construction of ANN potentials [110–114], and is available in the free and open-source ANN potential package `anet` [50].

The original high-dimensional ANN potential by Behler and Parrinello (BP) introduced a set of *symmetry functions* for the sampling of *bond lengths* and *bond angles* within local atomic environments [23, 99], which were in the spirit of earlier techniques for the symmetry-invariant representation of atomic coordinates [22]. Artrith, Urban, and Ceder (AUC) showed that the symmetry-function descriptor can be understood as a basis set expansion of the RDF and ADF, and proposed to replace the original BP functions with an orthonormal basis set of Chebyshev polynomials [100]. The expansion of the RDF of the local atomic environment of atom i can then be written as (the ADF expansion is analogous)

$$\text{RDF}_i(R) \approx \sum_{\alpha=0}^{\alpha_{\max}^{\text{RDF}}} c_\alpha^{\text{RDF}} \phi_\alpha(R) \quad \text{with} \quad c_\alpha^{\text{RDF}} = \sum_{\vec{R}_j \in \sigma_i} \bar{\phi}_\alpha(R_{ij}) f_c(R_{ij}), \quad (5)$$

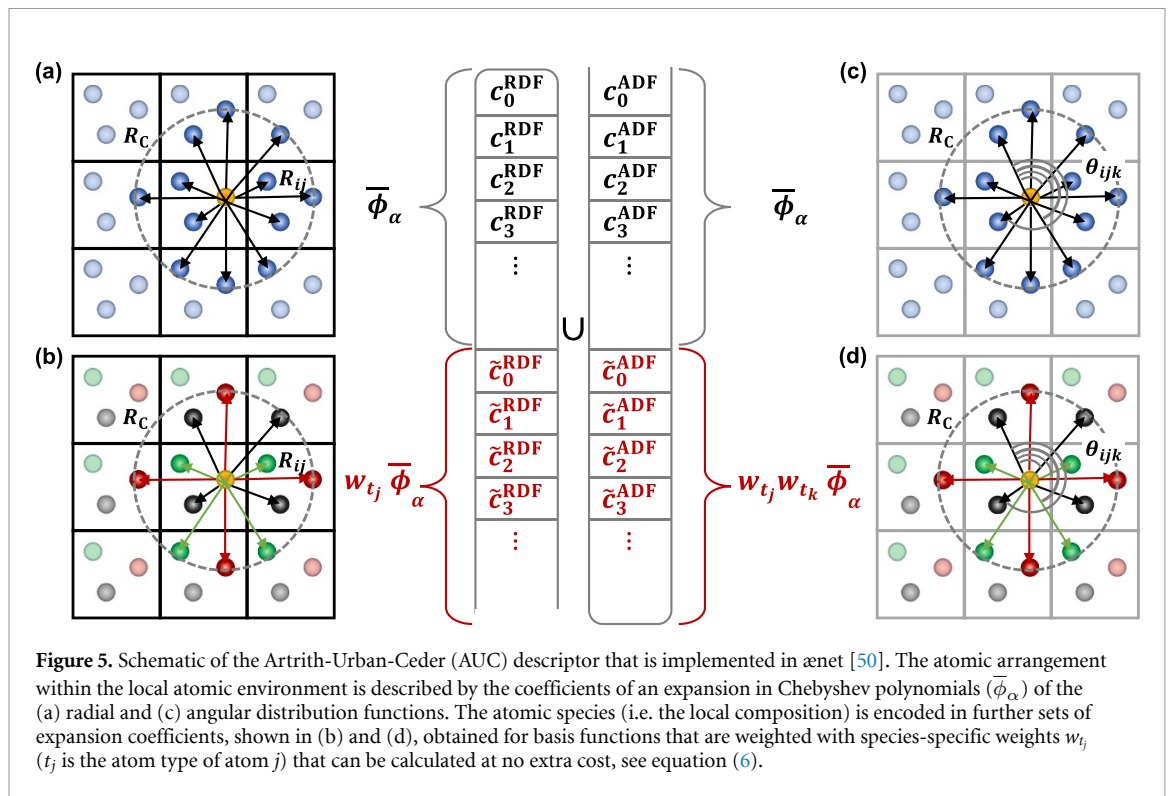
where $\bar{\phi}_\alpha$ is the Chebyshev polynomial of order α and ϕ_α is the orthonormal *dual* function. The cutoff function $f_c(R) = 0.5[\cos(\pi R/R_c) + 1]$ for $R \leq R_c$ goes smoothly to zero at the cutoff radius R_c and $R_{ij} = |\vec{R}_j - \vec{R}_i|$ is the distance of neighbor atom j from the central atom i . The RDF and ADF themselves already exhibit the symmetries of the potential energy, so that the joint sets of expansion coefficients $\{c_\alpha^{\text{RDF}}\} \cup \{c_\alpha^{\text{ADF}}\}$ can be used as an invariant feature vector, and the precision and the length of the AUC feature vector can be adjusted by deciding on the polynomial orders $\alpha_{\max}^{\text{RDF}}$ and $\alpha_{\max}^{\text{ADF}}$ at which the expansions are truncated.

Figure 4 shows the Chebyshev basis functions, and figures 5(a) and (c) show a schematic of the construction of the feature vector for the local atomic *structure*.

The expansion in equation (5) provides the featurization of the atomic positions within the local atomic environment, but it does not account for the different chemical species, since all atoms are considered equal. To encode also information about atom types, an additional set of expansion coefficients is included as features, for which the contributions from each atom are *weighted* with species-specific weights w_{t_j} for atom type t_j [100]

$$\begin{aligned} \tilde{c}_\alpha^{\text{RDF}} &= \sum_{\vec{R}_j \in \sigma_i} w_{t_j} \bar{\phi}_\alpha(R_{ij}) f_c(R_{ij}) \\ \tilde{c}_\alpha^{\text{ADF}} &= \sum_{\vec{R}_j \in \sigma_i} \sum_{\vec{R}_k \in \sigma_i} w_{t_j} w_{t_k} \bar{\phi}_\alpha(\theta_{ijk}) f_c(R_{ij}) f_c(R_{ik}), \end{aligned} \quad (6)$$

where θ_{ijk} is the cosine of the angle between atoms i , j , and k . This weighted expansion gives rise to two more sets of expansion coefficients $\{\tilde{c}_\alpha^{\text{RDF}}\}$ and $\{\tilde{c}_\alpha^{\text{ADF}}\}$ that can be calculated at essentially no additional cost together with the unweighted coefficients. This featurization of the *composition* is shown schematically in figures 5(b) and (d). Note that the length of the feature vector does not depend on the number of chemical species that are present in the atomic structure.



3.1.2. Selecting the descriptor resolution

The representation of the local atomic environment by expansion-based descriptors, such as the AUC descriptor discussed above, becomes more precise as the number of basis functions is increased. However, the computational cost of both the featurization and the evaluation of the atomic energy ANNs also depends on the number of features. The descriptor resolution should therefore be considered a hyperparameter that has to be optimized for a given application, so that the number of features is as large as needed and as small as possible.

For descriptors that are not based on expansion, an additional feature selection step can be introduced. *Feature engineering* and *feature selection* are subject of current research, and various methods have been proposed for constructing descriptors and selecting relevant features automatically [115–118].

3.1.3. Selecting a radial cutoff

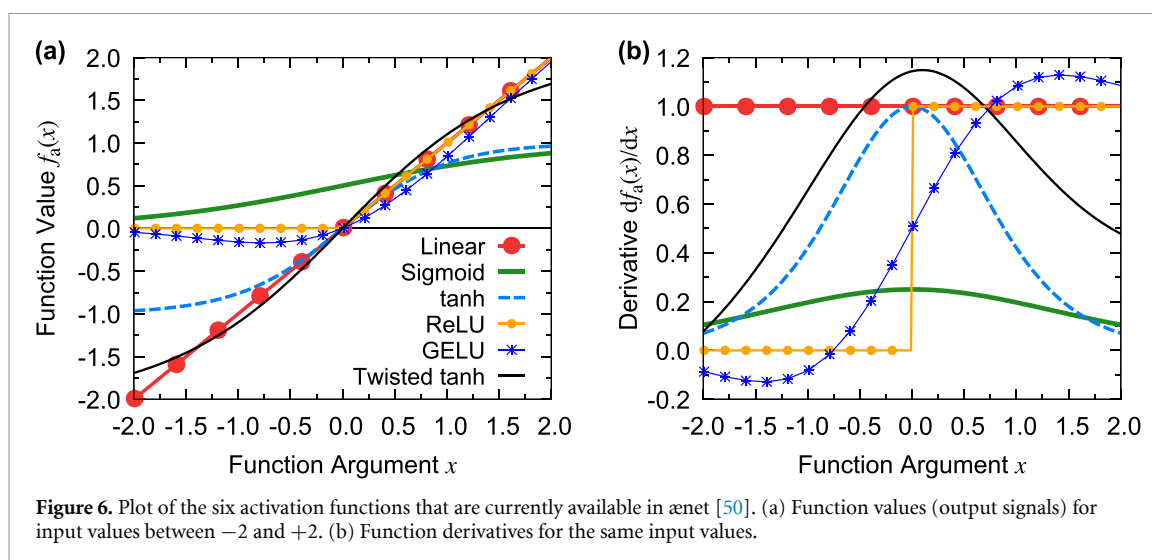
Another hyperparameter that should be optimized during the model construction phase is the cutoff radius R_c of the local atomic environments.

The larger the cutoff radius is, the more information is available for featurization, and the more atomic structures can, in principle, be distinguished. However, the computational cost of featurization also generally increases with the number of atoms within the local atomic environment, and the number of atoms scales as R_c^3 . For computational efficiency it is therefore beneficial to choose a radial cutoff that is as small as possible.

The radial cutoff R_c can also strongly affect the convergence of training. If R_c is chosen too small, the feature vectors might contain insufficient information for the construction of accurate ANN potentials, resulting in poor generalization. But if R_c is chosen too large, the feature vectors might become dominated by irrelevant structural and compositional differences that do not actually affect the atomic energies, thus leading to poor training convergence.

Typical cutoff radii lie after the second or third coordination shell of the central atom, which corresponds to $\sim 6\text{--}8$ Å for metal oxides [50, 100, 119, 120] and $\sim 4\text{--}6$ Å for organic molecules [51]. It can be beneficial to increase R_c further to capture also non-bonded interactions, such as hydrogen bonds, and for water cutoff distances of $6\text{--}10$ Å have been reported to give accurate results [12, 88, 121].

However, since the optimal cutoff range can be strongly dependent on the chemical system and application, it is necessary to perform a parameter study and compare the accuracy and transferability of the resulting ANNs in the validation stage of model construction (section 4).



3.2. Recipe: ANN model parameters

While the featurization determines the input dimension of the atomic energy ANNs, the internal architecture of the ANNs and the employed activation functions are also hyperparameters that need to be optimized.

3.2.1. Selecting the ANN architecture

The *architecture* of an atomic ANN is defined by the number of hidden layers and the number of neurons (nodes) in each hidden layer. As discussed in section 1.1, the number of nodes in the input layer is defined by the dimension of the feature vector, and the output layer consists of a single node that returns the atomic energy.

The ANN architecture thus determines the model *complexity*, in the sense that it determines the number of optimization parameters, i.e. the weights $\{w_{k,j}^i\}$ and $\{b_{ij}\}$ of equation (1). If the ANN is too small, i.e. if it has too few hidden layers or nodes per layer, the MLP will not be flexible enough to reproduce the underlying PES well. However, if the ANN is too large and too flexible, it might learn spurious information such as noise during training, which leads to overfitting (see section 4).

It is possible to monitor for and to avoid overfitting even with large ANN architectures, such as deep ANNs with more than three layers, but larger-than-needed ANN architectures are also undesirable because the computational effort for training and ANN potential evaluation increases approximately as N^2 with the number of neurons per layer N . Thus, generally the smallest ANN architecture yielding the required accuracy and transferability standards should be employed for constructing ANN potentials [122].

Many reported accurate ANN potentials are based on architectures with two hidden layers. The number of neurons per layer is a hyperparameter that has to be optimized for a given chemical system and application, though often ~ 20 – 30 nodes per layer can already yield highly accurate ANN potentials even for complex materials [31]. For large reference data sets and complex structure/composition spaces, architectures with more than two hidden layers and more than 50 nodes per layer may be beneficial [51].

3.2.2. Selecting the activation functions

The choice of the activation function f_a in equation (1) is another hyperparameter. For ANN training with standard backpropagation methods (see section 4), the activation function needs to be differentiable. The activation function needs to be non-linear, or otherwise the ANN function becomes a linear model. Further, it was found that monotonically increasing activation functions can accelerate the convergence of the weights during training [123].

Some of the most common activation functions and their derivatives used for atomic-energy ANNs are shown in figure 6. The activation potential of biological neurons resembles a step function, and step-like or sigmoidal activation functions, such as the logistic function or the hyperbolic tangent, are also a popular choice for artificial neurons. However, both functions are non-constant for only a small range of input values (activations), so that care must be taken during training to ensure a non-vanishing gradient of the *loss function* (section 4.1.3). To avoid such saturation issues, a constant linear slope can be added to the sigmoidal function, an example of which is the *twisted* hyperbolic tangent function [74] shown in figure 6. The rectified linear unit (ReLU) activation function has been introduced more recently [124] to avoid vanishing gradients in deep ANNs. Although the derivative of the ReLU function exhibits a discontinuity at

0, good training results can be obtained in practice also for regression models such as ANN potentials [125, 126]. The Gaussian error linear unit (GELU) function [127] has similar properties to ReLU but does not exhibit any discontinuity in its derivative and is thus an even better choice for ANN potentials. Many more activation functions have been proposed in the literature, and the development of activation functions is still an active area of research.

In conclusion, the activation functions for the hidden layers of ANN potentials have undergone an evolution over time, and the twisted hyperbolic tangent and the GELU function address the shortcomings of the previous generation and are typically good choices in practice.

Note that the range of output values that an artificial neuron can produce is determined by the co-domain of the activation function. Therefore, the activation function of the final layer of atomic-energy ANNs is typically chosen to be the linear function, so that the output of the ANNs, i.e. the atomic energy, is unconstrained.

3.3. Example: model selection

For the water example of section 2.2, the hyperparameters were optimized during model construction. The best compromise of accuracy on the validation set and computational efficiency was obtained for a model with a descriptor dimension of 51 (for hydrogen atoms) and 46 (for oxygen atoms), a radial cutoff of 6.35 Å, two hidden layers, and 25 nodes per layer with hyperbolic tangent activation function [88].

An example of an ANN potential for an inorganic solid material is the potential for amorphous LiSi alloys from [128], for which the hyperparameters were also thoroughly optimized. The final ANN potential employed an AUC descriptor with radial and angular expansion order of 10 (= 44 dimensions in total including the zeroth order), a cutoff radius of 8 Å, two hidden layers with each 15 nodes, and hyperbolic tangent activation function.

4. Model training and validation

For a specific choice of hyperparameters (section 3), the ANN potential model needs to be *trained* and *validated* on different parts of the data set (section 2). Model training is the process of optimizing the weight parameters $\{w_{k,j}^i\}$ and $\{b_{i,j}\}$ of equation (1) for all nodes of the ANN such that a *loss* function \mathcal{L} is minimized. Since the optimization targets, i.e. the total energies of the reference structures, are given, this is a *supervised learning* task. Combining all weight parameters in a single set W , the weight optimization can be expressed as

$$W_{\text{opt}} = \arg \min_W \mathcal{L}(W; S_{\text{train}}), \quad (7)$$

where W_{opt} is the set of optimal weight parameters and S_{train} is the *training set* of all atomic structures from the reference data set that are used for training. The remaining atomic structures make up the *validation set* S_{val} that is used during training to monitor the progress, to detect overfitting, and to obtain an initial estimate of the ANN potential accuracy. As indicated in the flow chart of figure 1, the model hyperparameters are typically varied until the model achieves optimal performance on the validation set.

4.1. Recipe

To perform the optimization of equation (7) in practice:

- (a) the data set needs to be split into training and validation sets,
- (b) the initial values for the weight parameters need to be set,
- (c) a suitable loss function needs to be defined, and
- (d) a training method has to be chosen.

4.1.1. Selecting training/validation sets

Both training set S_{train} and validation set S_{val} are derived from the overall reference data set. The training : validation split (point 1. of the above list) is often between 90% : 10% and 50% : 50%, depending on the size of the reference data set. The validation data should be selected randomly and should be representative for the entire reference data set. The training and validation sets must not overlap, i.e. no atomic structure may be present within both sets.

The validation set is used only for obtaining an initial estimate of the ANN potential accuracy and its ability to generalize, but additional independent testing of the trained model is necessary (section 5). The main purpose of the validation set is to find the optimal set of hyperparameters that minimizes the validation error on unseen structures.

4.1.2. Weight initialization and feature/target standardization

The accuracy that a trained ANN model can achieve and the efficiency of solving the weight optimization of equation (7) can strongly depend on the initial values of the weight parameters W as well as the value range of the features and targets [129]. Feature/target normalization and the choice of initial weight parameters is therefore an important first step for the training of ANN potentials.

As discussed in section 3.2.2 and shown in figure 6(b), the gradient of typical activation functions is non-zero only for a narrow range of input values, which in turn depend on the value of the weights $w_{k,j}^i$ and $b_{i,j}$ and the magnitude of the features or node outputs $x_{i,j}$ (equation (1)). If the initial weight parameters are chosen such that the activation function gradients are close to or identical to zero, standard methods for the weight optimization equation (7) that follow the gradient of the loss function are inefficient. This issue of vanishing gradients can be alleviated by ensuring that the output values of all neurons in each given layer are initially centered around zero [74].

A common approach is shifting and scaling the features (i.e. the descriptors of the local atomic environment of section 3.1) such that their values are centered around the point of inflection of the activation function (if applicable) and fall into the non-constant range of the activation function. For example, for hyperbolic tangent activation functions (see section 3.2.2), the features can be shifted and scaled such that their variance is equal to 1 and the values are zero-centered [74]. Note that this standardization should be applied to each feature individually. With this convention for feature standardization, the weight parameters should also be initialized such that the distribution of weights has a variance and center that is appropriate for the activation function used (e.g. a variance of 1 and center of 0 for the hyperbolic tangent). In practice, initial weights are first drawn from a random distribution and then normalized.

The convergence of the learning process generally benefits from adjusting also the weights of the hidden layers such that the arguments of their nodes reside in the non-linear region of the activation function. Various weight initialization methods have been developed and proposed in the literature, and an overview on common techniques can be found in [129] and the references therein.

Although ANN potentials typically use a linear activation function in the output layer so that the energy is unconstrained, it is still beneficial to scale the target values so that large differences in the scale of the weights in the output layer to those in the other layers are avoided, and the initial ANN output is already close to the target values. Thus, the structural energies are also commonly shifted and scaled such that they are zero-centered and have the same variance as the features. Instead of standardizing the target values, the outer ANN weights can also be adjusted before model training so that the mean and standard deviation of the initial ANN output matches the target distribution (see figure 7).

The impact of feature/target standardization and weight initialization on the initial state of an ANN potential is visualized in figure 7. Panel (a) of the figure shows that the initial predictions of the ANN potential before training can be orders of magnitude different from the DFT reference values, if the target energies and forces are not standardized. After proper standardization, the initial predictions are of the same order of magnitude as the target values as shown in figure 7(b), which generally accelerates the training process significantly.

4.1.3. Selecting a loss function

The loss function \mathcal{L} (sometimes also called *objective function*) encodes the optimization targets, i.e. the properties that the ANN potential is trained to reproduce, such as the total structural energy and/or the interatomic forces. Whether only energies or only forces should be included in the loss function depends on the planned application of the ANN potential, since the improved performance of the chosen target property usually causes the property that is left out to be of reduced accuracy. On the other hand, training on both energies and forces simultaneously is computationally more demanding than training on each quantity separately.

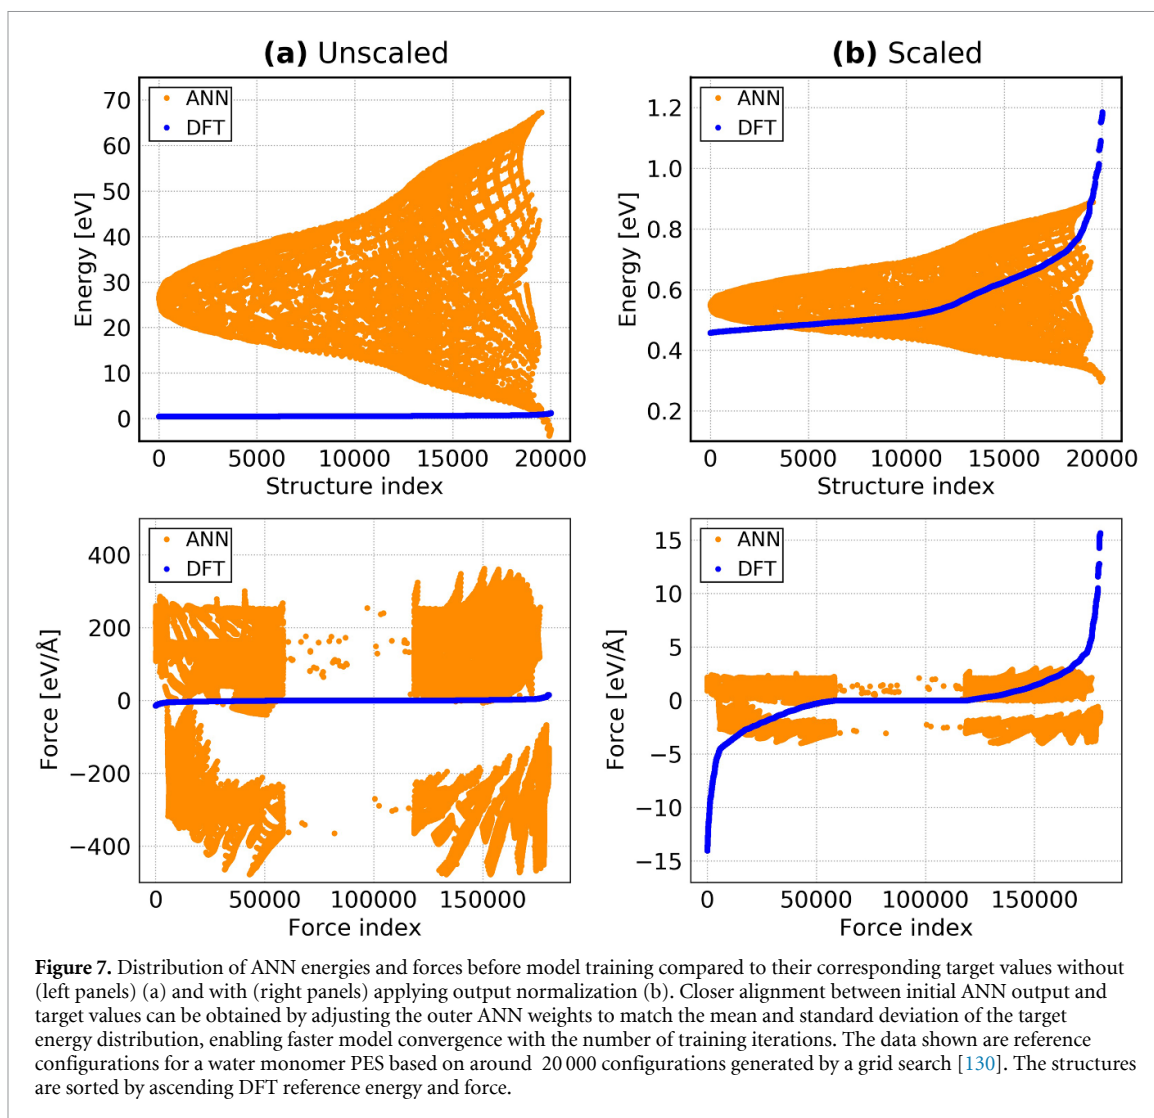
4.1.3.1. Energy training

The most common choice of loss function \mathcal{L} is for training on energy information only

$$\mathcal{L}_E(W) = \frac{1}{N_{\text{train}}} \sum_{\sigma \in S_{\text{train}}} \frac{1}{2} \left(E_{\text{ANN}}(\sigma; W) - E_{\text{ref}}(\sigma) \right)^2, \quad (8)$$

where N_{train} is the number of atomic structures in the training set S_{train} , E_{ANN} is the ANN potential energy of equation (3), E_{ref} is the corresponding reference energy, and the sum runs over all structures σ within the training set.

Although some simulation techniques do not require gradients, such as Metropolis Monte-Carlo sampling [131], for most applications including geometry optimizations and MD simulations the



interatomic forces need to be well represented. Training on the energy-dependent loss function \mathcal{L}_E of equation (8) can yield accurate forces, but a fine sampling of the phase space of interest might be required [132]. Hence, a training set with a large number of structures might be needed, and the computational cost for the reference calculations can become limiting. It is therefore often desirable to include the atomic forces as additional training targets [133].

4.1.3.2. Force training

In principle, it is possible to train an ANN potential on force information only by employing a loss function that is based on the force errors

$$\mathcal{L}_F(W) = \frac{1}{N_{\text{train}}} \sum_{\sigma \in S_{\text{train}}} \sum_i^{\text{atoms}} \left| \vec{F}_{i,\text{ANN}}(\sigma; W) - \vec{F}_{i,\text{ref}}(\sigma) \right|^2, \quad (9)$$

where $\vec{F}_{i,\text{ANN}}(\sigma; W)$ is the force acting on the i th atom in structure σ as predicted by the ANN potential and $\vec{F}_{i,\text{ref}}$ is the corresponding reference force. For efficient force training, the ANN potential can be expressed such that the final layer returns an atomic force vector instead of an atomic energy [134]. However, the total structural energy can only be obtained up to an unknown constant from ANN potentials that predict forces, which makes it challenging to validate such potentials for the prediction of thermodynamic quantities.

4.1.3.3. Energy and force training

A more comprehensive loss function can be constructed by combining the energy and force loss functions from equations (8) and (9), which can also be generalized to higher derivatives of the energy

$$\mathcal{L}_{E,F}(W) = a_E \mathcal{L}_E(W) + a_F \mathcal{L}_F(W) \quad (10)$$

where a_E and a_F are weights that determine the contributions of the energy and force errors to the overall loss function value.

The combined energy and force training ensures that information from both the potential energy and its gradient enter the training, which reduces the number of reference data points required for training [112]. However, since the forces are the negative gradient of the energy, training with the loss function of equation (10) requires the derivative of the Force with respect to the ANN weights (i.e. the second derivatives of the energy), which can become computationally demanding and is error-prone to implement in computer code. We note that complex implementations can be avoided by utilizing efficient numerical schemes instead of fully analytical derivatives [135].

Another alternative is the translation of force information to additional energy reference data points by approximating the energy of atomic structures with displaced coordinates using a Taylor expansion [112]. This approach improves the force prediction accuracy by increasing the density of training points without requiring additional reference calculations.

4.1.4. Training methods

Various methods for ANN weight optimization (equation (7)) have been proposed in the literature. Most practical methods make use of the gradient of the loss function, i.e. the derivative of the loss function \mathcal{L} with respect to the ANN weight parameters W , which can be efficiently calculated using error backpropagation [74, 136]. The choice of training method ultimately depends on the size of the training set, the size of the structures in the training set (in terms of atoms), and the available computer hardware.

In general, two classes of training methods are distinguished, *batch training* and *online training* methods. In batch training, the ANN weights are updated based on the actual value and gradient of the loss function, which requires evaluating the errors of all samples in the training set. In contrast, in online training, ANN weights are updated sequentially based on the errors and gradients of each individual sample from the reference data set. An intermediate approach is the online training with *mini batches*, in which the weight updates are calculated based on blocks of data containing a specified number of samples.

An advantage of batch-training methods is that the second weight derivative (the Hessian) can be approximated more readily, such as in the limited-memory Broyden-Fletcher-Goldfarb-Shanno (BFGS) method [137–141], which can enable convergence to solutions with lower residual loss function. Batch training can also be parallelized trivially, since the errors of all atomic structures in the reference data set can be evaluated simultaneously.

One advantage of online training with stochastic gradient descent [74] or related methods such as ADAM [142] is the efficient evaluation of weight updates. In addition, models trained with online methods often generalize well because of the regularizing effect of the approximate loss function evaluation [74]. In online training, the training *schedule* can be controlled, i.e. the order in which samples are chosen from the reference data set, which can be useful for preferential training of atomic structures with high errors or with low energies (Boltzmann weighting). More complex online-training methods, such as the extended Kalman filter [143, 144], can also make use of approximate Hessian information.

4.1.5. Overfitting and extrapolation

ANN potential training using the methods of the previous section minimizes the loss function of equation (7) for the training set only. However, good performance for the training data does not necessarily imply that the potential will generalize to structures that were not included in the training process.

In order to estimate how well the MLP generalizes, the loss obtained for the validation set (section 4.1.1) is typically considered. If the validation-set loss is similar to the training-set loss, it can be assumed that the ANN potential generalizes well for structures that are reasonably similar to those in the reference data set. If the validation-set loss is significantly larger than the training-set loss, either *overfitting* has occurred or the reference data set samples the structural space too sparsely so that *extrapolation* is observed.

Overfitting is the phenomenon when the ANN function reproduces the training samples with high accuracy but at the cost of introducing unreasonable behavior in the regions between the training points (see figure 8(a)). In the case of noisy reference data, overfitting also means that the ANN incorrectly reproduces the noise and not only the expected value of the targets. Overfitting occurs when the ANN is too flexible for the size and density of the training set, i.e. if the training set contains too few data points or if the model complexity is too great and an adjustment of the hyperparameters (see section 3) is needed. Overfitting can be reduced by introducing regularization terms in the loss function [74, 145] or by extending the training set, for example, by including force information (see figure 8(b)). During training, overfitting can be detected by monitoring the training-set and validation-set loss, which start to diverge at the onset of overfitting.

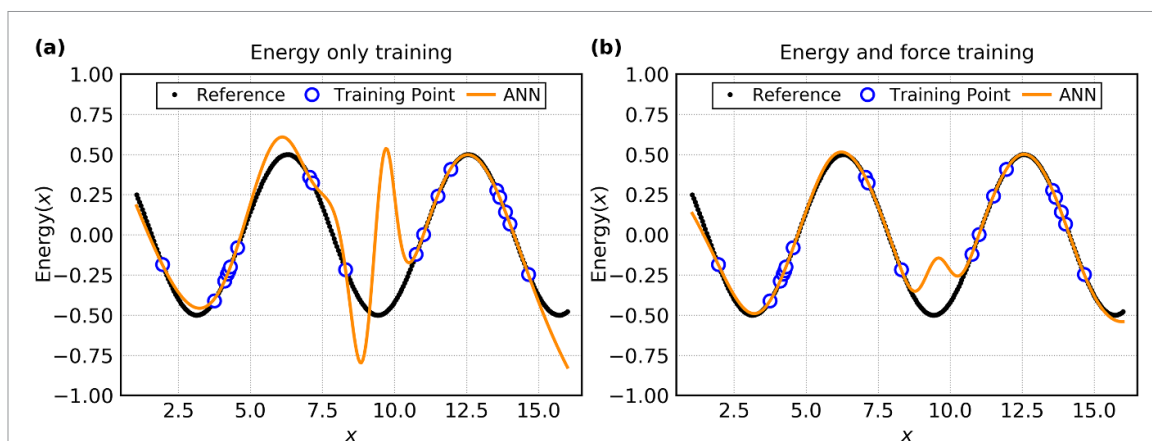


Figure 8. Overfitting and use of gradient information: Fitting of the simple model function $\cos(x)$ in order to illustrate the effect of overfitting. Here, x would correspond to an interatomic distance and the cosine function represents minima and maxima (or saddle points) of the PES. When only energy information is used to optimize the ANN weights, the training points are accurately interpolated but poor results for points not in the training set are found (a). Training the ANN to forces in addition to the reference energies, results in an improved representation of points not in the training set (b). For both examples, a $1 - 30 - 30 - 1$ ANN architecture was used, and the x value was the only input.

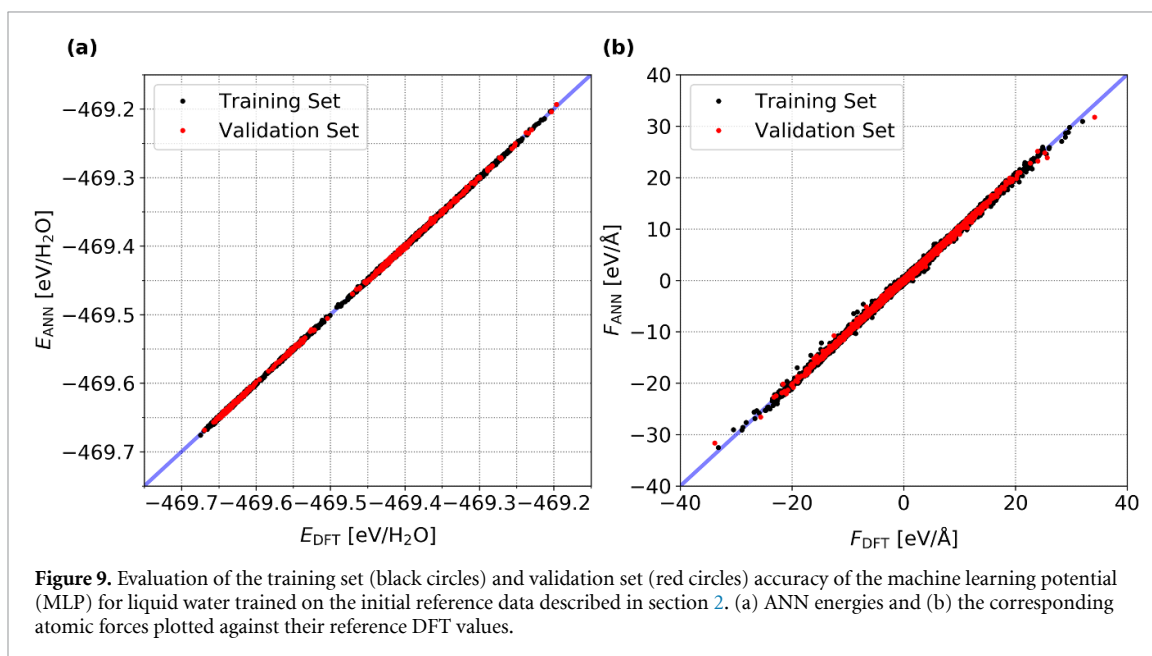


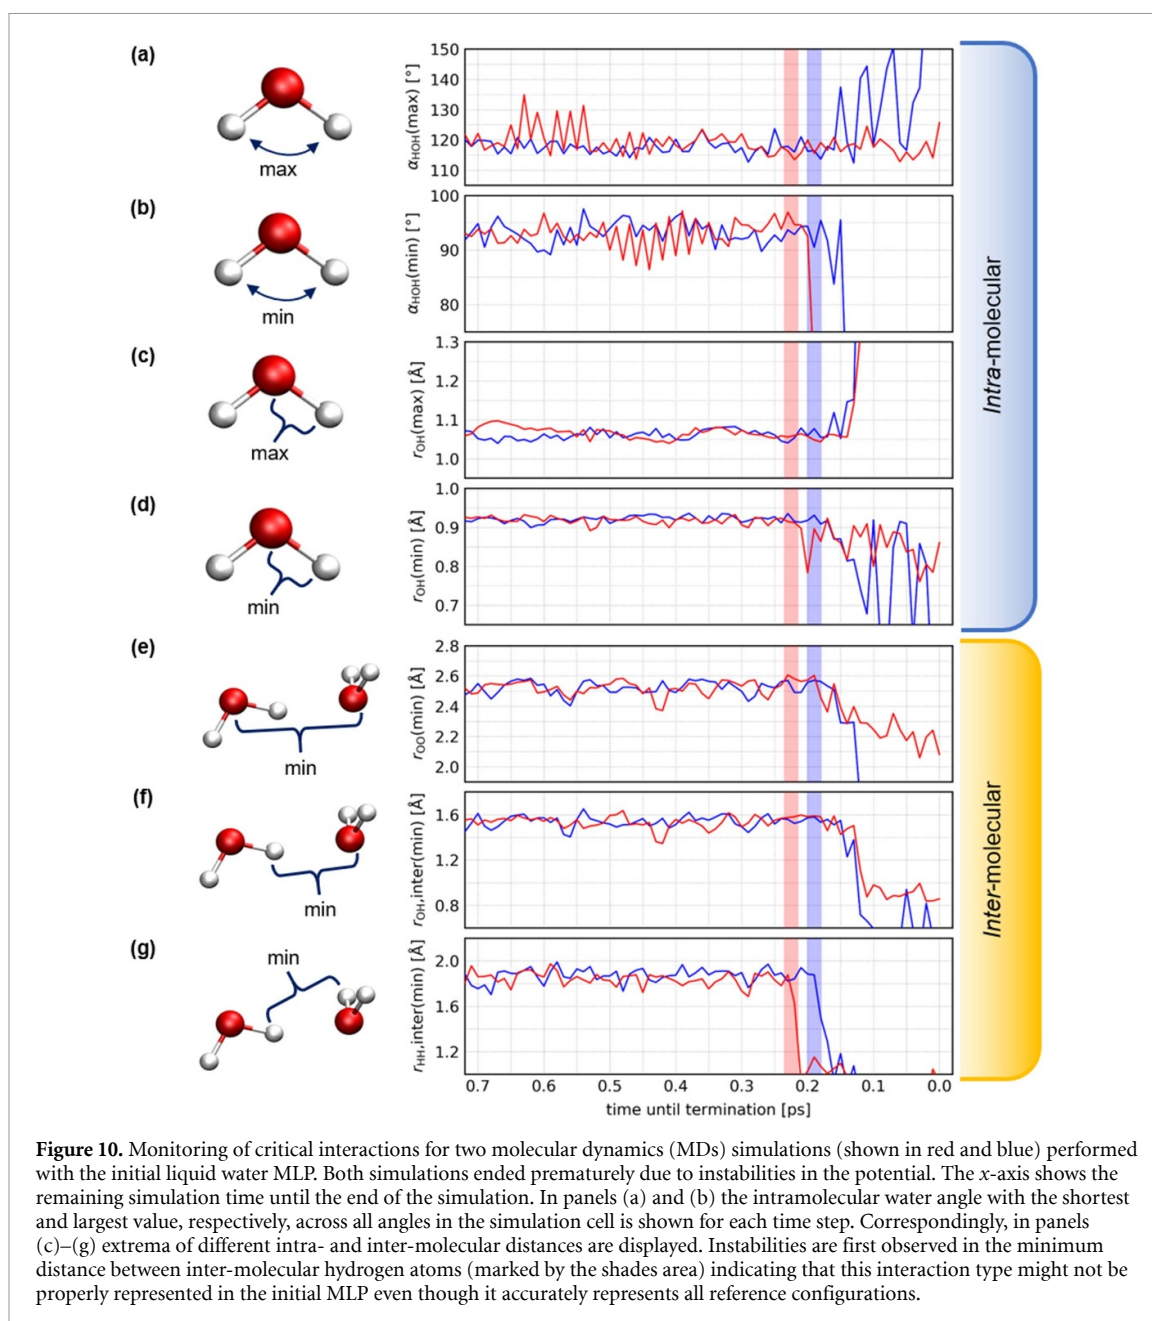
Figure 9. Evaluation of the training set (black circles) and validation set (red circles) accuracy of the machine learning potential (MLP) for liquid water trained on the initial reference data described in section 2. (a) ANN energies and (b) the corresponding atomic forces plotted against their reference DFT values.

As discussed in the data selection section 2, ANNs are unreliable for extrapolation [78]. If the training set samples the structure space too coarsely, regions of the PES may be insufficiently represented. An example is shown in figure 8(b). Underrepresented regions can also be identified by the validation set, if at least some relevant structures are present.

4.2. Example: training and validation

Returning to our example case on liquid water, figure 9 examines the training and validation set accuracy of the MLP trained on the initial data set described in section 2.1 based on a training : validation split of 90% : 10%. It can be seen that both training and validation set are accurately represented by the initial MLP across the full range of total energy and atomic force values with low validation root mean squared errors (RMSEs) (1.4 meV/H₂O and 83.2 meVÅ, respectively) that are on the same level as the corresponding training values (1.2 meV/H₂O and 84.7 meV Å⁻¹).

These low figures and the close correspondence between training and validation errors indicate that the initial MLP has been sufficiently optimized and a close to optimal choice of hyperparameters has been found. There are no obvious outliers and the accuracy is similar over the entire energy range, i.e. no energy region is underrepresented. Now, the model can be tested in the intended application to quantify its accuracy.



5. Model testing and active learning

Once an ANN potential has been obtained that performs well on the validation set (section 4), the potential needs to be tested in actual applications, such as MD simulations. Generally, an initial data set generated as described in section 2, e.g. by sampling from MD simulations and/or by manually perturbing equilibrium structures, does not ensure that the configuration space visited during the intended application is fully covered and can therefore lead to ANN potentials that exhibit instabilities in simulations. This is exemplified in our water case study when the initial MLP was first employed in MD simulations. As demonstrated in figure 10, the initial MLP exhibits stability issues leading to a premature termination of the simulation even though its validation accuracy is high (figure 9) and its initial training set comprises of a diverse set of structures which cover a broad energy and force range (figure 3). A detailed analysis of the corresponding trajectories reveals the underlying reason for the stability issue: the interactions between inter-molecular hydrogen atoms is not properly described.

Such generalization issues can be addressed by iteratively including additional data points in the reference data set, as shown in the outer loop of figure 1. Here, a model trained on an initial data set is used in preliminary applications and then *retrained* once it encounters configurations for which the model

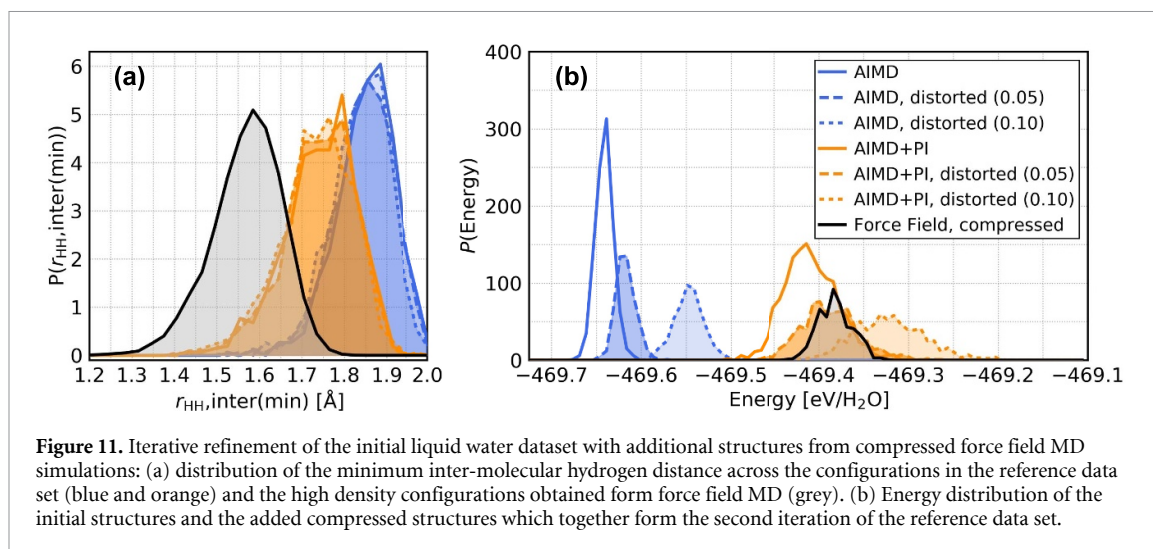


Figure 11. Iterative refinement of the initial liquid water dataset with additional structures from compressed force field MD simulations: (a) distribution of the minimum inter-molecular hydrogen distance across the configurations in the reference data set (blue and orange) and the high density configurations obtained from force field MD (grey). (b) Energy distribution of the initial structures and the added compressed structures which together form the second iteration of the reference data set.

prediction shows a low accuracy. The challenge here is how to identify such configurations with a high model uncertainty.

While in principle one could recompute all configurations generated with a given model by the underlying reference method to detect inaccurately described structures, this is computationally very inefficient. A solution to this is to make use of *active learning* approaches that select the next training set iteration from *unlabeled data* in an automated fashion with the benefit that the amount of expensive reference calculations is limited.

Active learning approaches for the construction of MLPs [122, 146, 147] comprise of three steps which are executed in a loop until the desired model performance is reached:

- (a) efficient *exploration* of the configuration space,
- (b) *selection* of relevant configurations and labeling, i.e. calculations of reference energies and forces,
- (c) followed by model *retraining*.

One can distinguish between on-the-fly active learning [147–149] in which retraining happens *during* the simulation and offline active learning where the next iteration of training structures is first accumulated, then a new model is trained, and subsequently a new simulation with the improved MLP will be launched.

The crucial step in an active learning approach is the selection of data points with high model uncertainty without knowing their reference properties beforehand. The goal is to find a query strategy to decide if a given configuration is already well described by the current ANN potential or if it should be added to the reference data set. Approaches proposed in the literature generally belong to one of the following classes:

(a) *Data set reduction approaches*: A large set of candidate configurations is generated with some sampling strategy, redundant/similar configurations are removed, and additional reference calculations are only performed for the configurations that are most different from those in the present reference data set. For example, Bernstein *et al* selected a subset of configurations from relaxation trajectories by (1) Leverage-score CUR algorithm, and (2) Flat histogram sampling (selection from low-density regions) with Boltzmann-probability bias [150].

(b) *Query-by-committee approaches*: An ensemble of models is used to evaluate a set of candidate configurations that were generated with some sampling strategy using one specific model. The standard deviation of the energy (and potentially the forces) across the committee of models is then used as an uncertainty estimate [79, 122, 151]. Ensemble methods can be also combined with data set reduction techniques as described above [79].

(c) *Statistical uncertainty quantification*: Statistical inference, e.g. based on the Bayesian framework [152], can be used to estimate the uncertainty of predictions. GPR models provide an intrinsic uncertainty estimate that can be employed [147, 149]. In ANN potentials, *dropout* can provide an uncertainty estimate [153], i.e. the *query-by-committee* can be built into the ANN architecture. In the case of MLP methods that depend linearly on the model parameters, such as moment tensor potentials [148], extrapolation can also be detected on-the-fly, which can be exploited for active learning.

5.1. Example: model testing and refinement

In the case of the liquid water MLP, the initial dataset was extended manually instead of using an active learning technique since the underlying reason for its stability issue could be identified by the analysis summarized in figure 10. To improve the representation of the ill-described inter-molecular hydrogen interactions additional reference structures were obtained from compressed MD simulations (employing an artificially increased density and a higher temperature) with a classical force field. As shown in figure 11, these structures sample short hydrogen-hydrogen distances and are located in a high-energy region. After labeling the additional structures with the corresponding DFT energies and forces, the retrained MLP could be applied in MD simulations without any stability issues even at elevated temperatures. Additional rounds of model testing and refinement at higher temperatures followed and the final reference set comprising of a total of 9189 structures can be obtained online from [154].

6. Discussion and outlook

In this tutorial review, we outlined common strategies for the construction of ANN-based machine-learning potentials. However, even if the best practices for the construction of MLPs are followed and active learning approaches are implemented for data selection, there are still remaining challenges that hinder the universal application of MLPs.

Depending on the size of the relevant configuration space, one significant challenge can be the expense of the reference calculations. Compiling a reference data set that captures all regions of the PES of a multi-component system can be a formidable challenge.

The construction of MLPs can be simplified by focusing on (a) the description of specific parts of the potential-energy, (b) certain regions of the system, or (c) using MLPs alongside full first-principles calculations, instead of describing the full PES of all atoms in the simulation box with a single MLP. Compared to general MLPs these *specialized MLP* approaches are generally more robust, require a lower number of expensive high-level reference calculations, and are easier to converge during training, while having the downside of being computationally more involved during the execution of the actual simulation.

Potentials fitted not on the full PE but on the differences to some reference are commonly referred to as *delta ML approaches*.

6.1. Energy decomposition approaches

Rather than learning the full potential energy, long-range contributions such as electrostatics or van der Waals (vdW) interactions can be removed before model training and the remaining short-range energy is used as the target property which can be more easily described by atomic ANNs that depend on local environments.

$$E^{\text{total}} = E_{\text{ANN}}^{\text{short}} + E_{\text{ANN}}^{\text{elec}} + E^{\text{vdW}}. \quad (11)$$

The removed energy contributions can then be added back in by employing expressions that explicitly consider the physical distance dependence. This can be done by either employing already available analytic expression, as in the case of vdW interactions [90, 155] (e.g. with Grimme's D2/D3 methods [156, 157]), or by training separate ML models, for example to represent atomic charges for calculating long-range electrostatics based on Coulomb's law [76, 121, 155]. A dependence of the MLP energy on long-ranged features can also be directly included, avoiding the need to explicitly model atomic charges which are not uniquely defined observables [81].

6.2. Energy-difference approaches

Another group of delta ML approaches focuses on the prediction of energy differences between two reference methods of different quality. Here the energy difference between a lower-level method and a higher-level quantum-mechanics (QMs) based method is predicted by an MLP. An early example of such a composite strategy in which an ML correction is added to a computationally efficient but less accurate QM method is the delta-machine learning approach by Ramakrishnan *et al* [158]. Other examples using different levels of theory are discussed in [111] and [158]. Instead of predicting *total energy differences*, molecular-orbital-based schemes model high-level electronic structure correlation energies using inputs from Hartree–Fock calculations with the goal of being more transferable across chemical systems [159, 160].

6.3. Embedding approaches

Following the spirit of QM/MM approaches [161–166], delta MLPs can be developed to focus on certain regions in space within the entire system. Those regions of the system that are not described by the MLP are instead described by molecular mechanics-based or lower-level QM methods [167]. This strategy can also be

combined with energy-difference approaches, for example by predicting the energy difference between a low-level QM method and a high-level QM method by an MLP for atoms inside a limited spatial region which is coupled to molecular mechanics-based interactions for describing the larger environment (making it a QM/MM/ML approach in which the *delta* refers to energy difference *and* spatial separation) [168].

6.4. Domain-limited approaches

Finally, *specialized* MLPs can be employed alongside conventional first principles calculations to accelerate the calculation of QM properties. These potentials are often focused on a specific region of the configurational space and generally do not need to be trained to the highest degree of accuracy since final first principles reference calculations are included as part of the full workflow. Examples are the use of ML-accelerated geometry optimizations in which initial structures are pre-optimized with an MLP, followed by a final *ab initio* optimization that requires fewer steps to convergence since its input structure is already close to the ground state [169, 170]. Specialized MLPs have also been employed in combination with evolutionary algorithms to determine the phase diagram of amorphous alloys [128].

7. Summary

Machine-learning interatomic potentials enable accurate and efficient atomic-scale simulations of complex systems that are not accessible by conventional first principles methods, but for many systems of interest machine-learning potentials have not yet been developed. Here, we reviewed common strategies and best practices for the construction of machine-learning potentials based on ANNs. Data selection, model selection, training/validation, and testing/refinement have been exemplified using practical examples. The number of available tools and data sets for machine-learning potential applications is rapidly growing, and we refer the reader to a curated and editable list at <https://github.com/atomisticnet/tools-and-data> with a collection of free and open-source tools and data resources. As discussed, the construction of ANN potentials is still a complex and manual process involving many steps. Recipes are provided here with the hope that in future more automated and standardized workflows for ANN construction will be established, so that the method can achieve its full potential in accelerating the prediction of materials and molecular properties with an unparalleled combination of accuracy and speed.

Data availability statement

The data that support the findings of this study are openly available at the following URL/DOI: <https://doi.org/10.24435/materialscloud:2020.0037/v1> and <https://doi.org/10.24435/materialscloud:dx-ct>.

Acknowledgments

N A acknowledges support by the Columbia Center for Computational Electrochemistry (CCCE). N A and A U thank the Extreme Science and Engineering Discovery Environment (XSEDE), which is supported by the National Science Foundation under Grant No. ACI-1053575, for supporting the development of the atomic energy network (*aenet*) package. A M M and J acknowledge support by the state of Baden-Württemberg through bwHPC and the German Research Foundation (DFG) through Grant No. INST 40/575-1 FUGG (JUSTUS 2 cluster). A M M and J K also thank the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) for supporting this work by funding—EXC2075—390740016 under Germany's Excellence Strategy and acknowledge the support by the Stuttgart Center for Simulation Science (SimTech). T M thanks T E Markland and O Marsalek for fruitful discussions.

Conflict of interest

The authors declare no competing financial interests.

Code details

Companion materials can be accessed at <https://github.com/atomisticnet/MLP-beginners-guide>, which includes a Jupyter notebook demonstrating the construction of an ANN potential and a second notebook that allows reproducing the failure analysis of figure 10.

ORCID iDs

April M Miksch  <https://orcid.org/0000-0002-3776-515X>
Tobias Morawietz  <https://orcid.org/0000-0002-9385-8721>
Johannes Kästner  <https://orcid.org/0000-0001-6178-7669>
Alexander Urban  <https://orcid.org/0000-0002-9021-279X>
Nongnuch Artrith  <https://orcid.org/0000-0003-1153-6583>

References

- [1] Hohenberg P and Kohn W 1964 *Phys. Rev.* **136** B864
- [2] Kohn W and Sham L J 1965 *Phys. Rev.* **140** A1133
- [3] Burke K 2012 *J. Chem. Phys.* **136** 150901
- [4] Becke A D 2014 *J. Chem. Phys.* **140** 18A301
- [5] Jones R O 2015 *Rev. Mod. Phys.* **87** 897
- [6] Mardirossian N and Head-Gordon M 2017 *Mol. Phys.* **115** 2315
- [7] Behler J, Martoňák R, Donadio D and Parrinello M 2008 *Phys. Status Solidi b* **245** 2618
- [8] Khaliullin R Z, Eshet H, Kühne T D, Behler J and Parrinello M 2010 *Phys. Rev. B* **81** 100103
- [9] Khaliullin R Z, Eshet H, Kühne T D, Behler J and Parrinello M 2011 *Nat. Mater.* **10** 693
- [10] Sosso G C, Donadio D, Caravati S, Behler J and Bernasconi M 2012 *Phys. Rev. B* **86** 104301
- [11] Natarajan S K and Behler J 2016 *Phys. Chem. Chem. Phys.* **18** 28704
- [12] Morawietz T, Singraber A, Dellago C and Behler J 2016 *Proc. Natl Acad. Sci.* **113** 8368
- [13] Rowe P, Csányi G, Alfè D and Michaelides A 2018 *Phys. Rev. B* **97** 054303
- [14] Stricker M, Yin B, Mak E and Curtin W A 2020 *Phys. Rev. Mater.* **4** 103602
- [15] Behler J 2016 *J. Chem. Phys.* **145** 170901
- [16] Artrith N 2019 *J. Phys. Energy* **1** 032002
- [17] Mueller T, Hernandez A and Wang C 2020 *J. Chem. Phys.* **152** 050902
- [18] Noé F, Tkatchenko A, Müller K-R and Clementi C 2020 *Annu. Rev. Phys. Chem.* **71** 361
- [19] Unke O T, Koner D, Patra S, Käser S and Meuwly M 2020 *Mach. Learn.: Sci. Technol.* **1** 0130001
- [20] Morawietz T and Artrith N 2021 *J. Comput. Aided Mol. Des.* **35** 557
- [21] Unke O T, Chmiela S, Sauceda H E, Gastegger M, Poltavsky I, Schütt K T, Tkatchenko A and Müller K-R 2021 *Chem. Rev.* accepted (<https://doi.org/10.1021/acs.chemrev.0c01111>)
- [22] Lorenz S, Groß A and Scheffler M 2004 *Chem. Phys. Lett.* **395** 210
- [23] Behler J and Parrinello M 2007 *Phys. Rev. Lett.* **98** 146401
- [24] Bartók A P, Payne M C, Kondor R and Csányi G 2010 *Phys. Rev. Lett.* **104** 136403
- [25] Møller C and Plesset M S 1934 *Phys. Rev.* **46** 618
- [26] Coester F and Kümmel H 1960 *Nuclear Phys. B* **17** 477
- [27] Čížek J 1966 *J. Chem. Phys.* **45** 4256
- [28] Bartlett R J and Musiał M 2007 *Rev. Mod. Phys.* **79** 291
- [29] Cremer D 2011 *WIREs Comput. Mol. Sci.* **1** 509
- [30] Zhang I Y and Grüneis A 2019 *Front. Mater.* **6** 123
- [31] Behler J 2017 *Angew. Chem., Int. Ed.* **56** 12828
- [32] Bartók A P, Kondor R and Csányi G 2013 *Phys. Rev. B* **87** 184115
- [33] Botu V and Ramprasad R 2015 *Int. J. Quantum Chem.* **115** 1074
- [34] Huan T D, Batra R, Chapman J, Krishnan S, Chen L and Ramprasad R 2017 *Comput. Mater.* **3** 37
- [35] John S T and Csányi G 2017 *J. Phys. Chem. B* **121** 10934
- [36] Nyshadham C, Rupp M, Bekker B, Shapeev A V, Mueller T, Rosenbrock C W, Csányi G, Wingate D W and Hart G L W 2019 *npj Comput. Mater.* **5** 51
- [37] Shapeev A V 2016 *Multiscale Model. Simul.* **14** 1153
- [38] Novikov I S, Gubaev K, Podryabinkin E V and Shapeev A V 2021 *Mach. Learn.: Sci. Technol.* **2** 025002
- [39] Gilmer J, Schoenholz S S, Riley P F, Vinyals O and Dahl G E 2017 *Proc. 34th Int. Conf. on Machine Learning (Proc. Machine Learning Research vol 70)* ed D Precup and Y W Teh (Sydney: PMLR, International Convention Centre) pp 1263–72
- [40] Duvenaud D K, Maclaurin D, Iparraguirre J, Bombarell R, Hirzel T, Aspuru-Guzik A and Adams R P 2015 *Advances in Neural Information Processing Systems 28* ed C Cortes, N D Lawrence, D D Lee, M Sugiyama and R Garnett (Curran Associates, Inc.) pp 2224–32
- [41] Kearnes S, McCloskey K, Berndl M, Pande V and Riley P 2016 *J. Comput. Aided Mol. Des.* **30** 595
- [42] Schütt K T, Arbabzadah F, Chmiela S, Müller K R and Tkatchenko A 2017 *Nat. Commun.* **8** 13890
- [43] Chen C, Ye W, Zuo Y, Zheng C and Ong S P 2018 (arXiv:1812.05055 [cond-mat, physics: physics])
- [44] Jørgensen P B, Jacobsen K W and Schmidt M N 2018 *Neural Message Passing with Edge Updates for Predicting Properties of Molecules and Materials* (arXiv:1806.03146)
- [45] Schütt K T, Sauceda H E, Kindermans P-J, Tkatchenko A and Müller K-R 2018 *J. Chem. Phys.* **148** 241722
- [46] Unke O T and Meuwly M 2019 *J. Chem. Theory Comput.* **15** 3678
- [47] Xie T and Grossman J C 2018 *Phys. Rev. Lett.* **120** 145301
- [48] Thompson A, Swiler L, Trott C, Foiles S and Tucker G 2015 *J. Comp. Phys.* **285** 316
- [49] Wood M A and Thompson A P 2018 *J. Chem. Phys.* **148** 241721
- [50] Artrith N and Urban A 2016 *Comput. Mater. Sci.* **114** 135
- [51] Smith J S, Isayev O and Roitberg A E 2017 *Chem. Sci.* **8** 3192
- [52] Zhang L, Han J, Wang H, Car R and Weinan E 2018 *Phys. Rev. Lett.* **120** 143001
- [53] Westermayr J, Gastegger M, Menger M F S J, Mai S, González L and Marquetand P 2019 *Chem. Sci.* **10** 8100
- [54] Tibshirani R 1996 *J. R. Stat. Soc. B* **58** 267

- [55] Hastie T, Tibshirani R and Friedman J H 2009 *The Elements of Statistical Learning: Data Mining, Inference and Prediction* 2nd edn (*Springer Series in Statistics*) (New York: Springer)
- [56] Mueller T and Ceder G 2009 *Phys. Rev. B* **80** 024103
- [57] Brown W M, Thompson A P and Schultz P A 2010 *J. Chem. Phys.* **132** 024108
- [58] Balabin R M and Lomakina E I 2011 *Phys. Chem. Chem. Phys.* **13** 11710
- [59] Hansen K, Montavon G, Biegler F, Fazli S, Rupp M, Scheffler M, von Lilienfeld O A, Tkatchenko A and Müller K-R 2013 *J. Chem. Theory Comput.* **9** 3404
- [60] Li Z, Kermode J R and De Vita A 2015 *Phys. Rev. Lett.* **114** 096405
- [61] Seko A, Takahashi A and Tanaka I 2015 *Phys. Rev. B* **92** 054113
- [62] Chmiela S, Tkatchenko A, Sauceda H E, Poltavsky I, Schütt K T and Müller K-R 2017 *Sci. Adv.* **3** e1603015
- [63] Gastegger M, Behler J and Marquetand P 2017 *Chem. Sci.* **8** 6924
- [64] Butler K T, Davies D W, Cartwright H, Isayev O and Walsh A 2018 *Nature* **559** 547
- [65] Cao L, Li C and Mueller T 2018 *J. Chem. Inf. Model.* **58** 2401
- [66] Hy T S, Trivedi S, Pan H, Anderson B M and Kondor R 2018 *J. Chem. Phys.* **148** 241745
- [67] Mardt A, Pasquali L, Wu H and Noé F 2018 *Nat. Commun.* **9** 5
- [68] Ryczko K, Mills K, Luchak I, Homenick C and Tamblyn I 2018 *Comput. Mater. Sci.* **149** 134
- [69] Artrith N 2020 *Matter* **3** 985
- [70] Artrith N, Lin Z and Chen J G 2020 *ACS Catal.* **10** 9438
- [71] Deringer V L, Caro M A and Csányi G 2019 *Adv. Mater.* **31** 1902765
- [72] Jørgensen W L and Tirado-Rives J 2005 *Proc. Natl Acad. Sci.* **102** 6665
- [73] Becker C A, Tavazza F, Trautt Z T and Buarque de Macedo R A 2013 *Curr. Opin. Solid State Mater. Sci.* **17** 277 *Frontiers in Methods for Materials Simulations*
- [74] Montavon G, Orr G B and Müller K-R (eds) 2012 *Neural Networks: Tricks of the Trade* 2nd edn (*Lecture Notes in Computer Science* vol 7700) (Berlin: Springer)
- [75] Blank T B, Brown S D, Calhoun A W and Doren D J 1995 *J. Chem. Phys.* **103** 4129
- [76] Artrith N, Morawietz T and Behler J 2011 *Phys. Rev. B* **83** 153101
- [77] Huang Y, Kang J, Goddard W A and Wang L-W 2019 *Phys. Rev. B* **99** 064103
- [78] Haley P and Soloway D 1992 [*Proc. 1992 IJCNN Int. Conf. on Neural Networks* vol 4 (Baltimore, MD: IEEE) pp 25–30]
- [79] Smith J S, Nebgen B, Lubbers N, Isayev O and Roitberg A E 2018 *J. Chem. Phys.* **148** 241733
- [80] Loeffler T D, Patra T K, Chan H, Cherukara M and Sankaranarayanan S K R S 2020 *J. Phys. Chem. C* **124** 4907
- [81] Grisafi A and Ceriotti M 2019 *J. Chem. Phys.* **151** 204105
- [82] Ghasemi S A, Hofstetter A, Saha S and Goedecker S 2015 *Phys. Rev. B* **92** 045131
- [83] Faraji S, Ghasemi S A, Rostami S, Rasoulkhani R, Schaefer B, Goedecker S and Amsler M 2017 *Phys. Rev. B* **95** 104105
- [84] Ko T W, Finkler J A, Goedecker S and Behler J 2020 *Nat. Commun.* **12** 398
- [85] Litman Y, Behler J and Rossi M 2020 *Faraday Discuss.* **221** 526
- [86] Cooper A M, Hallmen P P and Kästner J 2018 *J. Chem. Phys.* **148** 094106
- [87] Singraber A, Morawietz T, Behler J and Dellago C 2019 *J. Chem. Theory Comput.* **15** 3075
- [88] Morawietz T, Marsalek O, Pattenaude S R, Streacker L M, Ben-Amotz D and Markland T E 2018 *J. Phys. Chem. Lett.* **9** 851
- [89] Marsalek O and Markland T E 2017 *J. Phys. Chem. Lett.* **8** 1545
- [90] Morawietz T and Behler J 2013 *J. Phys. Chem. A* **117** 7356
- [91] Morawietz T, Urbina A S, Wise P K, Wu X, Lu W, Ben-Amotz D and Markland T E 2019 *J. Phys. Chem. Lett.* **10** 6067
- [92] Cheng B, Engel E A, Behler J, Dellago C and Ceriotti M 2019 *Proc. Natl Acad. Sci.* **116** 1110
- [93] Markland T E and Ceriotti M 2018 *Nat. Rev. Chem.* **2** 0109
- [94] Parsaefard B, De D S, Christensen A S, Faber F A, Kocer E, De S, Behler J, von Lilienfeld A and Goedecker S 2020 *Mach. Learn.: Sci. Technol.* **2** 015018
- [95] Khorshidi A and Peterson A A 2016 *Comput. Phys. Commun.* **207** 310
- [96] Unke O T and Meuwly M 2018 *J. Chem. Phys.* **148** 241708
- [97] Kocer E, Mason J K and Erturk H 2019 *J. Chem. Phys.* **150** 154102
- [98] Zaverkin V and Kästner J 2020 *J. Chem. Theory Comput.* **16** 5410
- [99] Behler J 2011 *J. Chem. Phys.* **134** 074106
- [100] Artrith N, Urban A and Ceder G 2017 *Phys. Rev. B* **96** 014112
- [101] Gastegger M, Schwiedrzik L, Bittermann M, Berzsenyi F and Marquetand P 2018 *J. Chem. Phys.* **148** 241709
- [102] Faber F A, Christensen A S, Huang B and von Lilienfeld O A 2018 *J. Chem. Phys.* **148** 241717
- [103] Christensen A S, Bratholm L A, Faber F A and Anatole von Lilienfeld O 2020 *J. Chem. Phys.* **152** 044107
- [104] Gubaev K, Podryabinkin E V and Shapeev A V 2018 *J. Chem. Phys.* **148** 241727
- [105] Reveil M and Clancy P 2018 *Mol. Syst. Des. Eng.* **3** 431
- [106] Himanen L, Jäger M O, Morooka E V, Federici Canova F, Ranawat Y S, Gao D Z, Rinke P and Foster A S 2020 *Comput. Phys. Commun.* **247** 106949
- [107] Lubbers N, Smith J S and Barros K 2018 *J. Chem. Phys.* **148** 241715
- [108] Schütt K T, Kessel P, Gastegger M, Nicoli K A, Tkatchenko A and Müller K-R 2019 *J. Chem. Theory Comput.* **15** 448
- [109] Nikitin F, Isayev O and Strijov V 2020 *Phys. Chem. Chem. Phys.* **22** 26478
- [110] Lacivita V, Artrith N and Ceder G 2018 *Chem. Mater.* **30** 7077
- [111] Sun G and Sautet P 2019 *J. Chem. Theory Comput.* **15** 5614
- [112] Cooper A M, Kästner J, Urban A and Artrith N 2020 *npj Comput. Mater.* **6** 54
- [113] Chen M S, Zuehlsdorff T J, Morawietz T, Isborn C M and Markland T E 2020 *J. Phys. Chem. Lett.* **11** 7559
- [114] Mori H and Ozaki T 2020 *Phys. Rev. Mater.* **4** 040601(R)
- [115] Cubuk E D, Malone B D, Onat B, Waterland A and Kaxiras E 2017 *J. Chem. Phys.* **147** 024104
- [116] Imbalzano G, Anelli A, Giofré D, Klees S, Behler J and Ceriotti M 2018 *J. Chem. Phys.* **148** 241730
- [117] Jinnouchi R, Karsai F, Verdi C, Asahi R and Kresse G 2020 *J. Chem. Phys.* **152** 234102
- [118] Li L, Li H, Seymour I D, Koziol L and Henkelman G 2020 *J. Chem. Phys.* **152** 224102
- [119] Artrith N, Hiller B and Behler J 2013 *Phys. Status Solidi b* **250** 1191
- [120] Elias J S, Artrith N, Bugnet M, Giordano L, Botton G A, Kolpak A M and Shao-Horn Y 2016 *ACS Catal.* **6** 1675
- [121] Morawietz T, Sharma V and Behler J 2012 *J. Chem. Phys.* **136** 064103

- [122] Artrith N and Behler J 2012 *Phys. Rev. B* **85** 045439
- [123] Wu H 2009 *Inf. Sci.* **179** 3432
- [124] Nair V and Hinton G E 2010 *ICML* pp 807–14
- [125] Han J, Zhang L, Car R and Weinan E 2018 *Commun. Comput. Phys.* **23** 629
- [126] Pattanaik P, Raghunathan S, Kalluri T, Bhimalapuram P, Jawahar C V and Priyakumar U D 2020 *J. Phys. Chem. A* **124** 6954
- [127] Hendrycks D and Gimpel K 2020 (arXiv:1606.08415[cs])
- [128] Artrith N, Urban A and Ceder G 2018 *J. Chem. Phys.* **148** 241711
- [129] Thimm G and Fiesler E 1997 *IEEE Trans. Neural Netw.* **8** 349
- [130] Morawietz T 2010 Entwicklung eines effizienten Potentials für das Wasser-Dimer basierend auf künstlichen neuronalen Netzen PhD Thesis school Master's thesis, Lehrstuhl für Theoretische Chemie, Ruhr-Universität Bochum
- [131] Artrith N and Kolpak A M 2014 *Nano Lett.* **14** 2670
- [132] Witkoskie J B and Doren D J 2005 *J. Chem. Theory Comput.* **1** 14
- [133] Christensen A S and von Lilienfeld O A 2020 *Mach. Learn.: Sci. Technol.* **1** 045018
- [134] Li W and Ando Y 2018 *Phys. Chem. Chem. Phys.* **20** 30006
- [135] Smith J S, Lubbers N, Thompson A P and Barros K 2020 (arXiv:2006.05475)
- [136] Rumelhart D E, Hinton G E and Williams R J 1986 *Nature* **323** 533
- [137] Broyden C G 1970 *IMA J. Appl. Math.* **6** 76
- [138] Fletcher R 1970 *Comput. J.* **13** 317
- [139] Goldfarb D 1970 *Math. Comp.* **24** 23
- [140] Shanno D F 1970 *Math. Comp.* **24** 647
- [141] Liu D C and Nocedal J 1989 *Math. Program.* **45** 503
- [142] Kingma D P and Ba J 2014 (arXiv:1412.6980 [cs.LG])
- [143] Blank T B and Brown S D 1994 *J. Chemom.* **8** 391
- [144] Julier S and Uhlmann J 2004 *Proc. IEEE* **92** 401
- [145] Hoerl A E and Kennard R W 1970 *Technometrics* **12** 55
- [146] Lookman T, Balachandran P V, Xue D and Yuan R 2019 *npj Comput. Mater.* **5** 1
- [147] Jinnouchi R, Miwa K, Karsai F, Kresse G and Asahi R 2020 *J. Phys. Chem. Lett.* **11** 6946
- [148] Podryabinkin E V and Shapeev A V 2017 *Comput. Mater. Sci.* **140** 171
- [149] Vandermause J, Torrisi S B, Batzner S, Xie Y, Sun L, Kolpak A M and Kozinsky B 2020 *npj Comput. Mater.* **6** 20
- [150] Bernstein N, Csányi G and Deringer V L 2019 *npj Comput. Mater.* **5** 99
- [151] Schran C, Brezina K and Marsalek O 2020 *J. Chem. Phys.* **153** 104105
- [152] George G C T and Box E P 2011 *Bayesian Inference in Statistical Analysis* (New York: Wiley)
- [153] Wen M and Tadmor E B 2020 *npj Comput. Mater.* **6** 124
- [154] Chen M S, Morawietz T, Markland T E and Artrith N 2020 AENET-LAMMPS and AENET-TINKER: interfaces for accurate and efficient molecular dynamics simulations with machine learning potentials *Materials Cloud Archive*
- [155] Yao K, Herr J E, Toth D W, Mckintyre R and Parkhill J 2018 *Chem. Sci.* **9** 2261
- [156] Grimme S 2006 *J. Comput. Chem.* **27** 1787
- [157] Grimme S, Antony J, Ehrlich S and Krieg H 2010 *J. Chem. Phys.* **132** 154104
- [158] Ramakrishnan R, Dral P O, Rupp M and von Lilienfeld O A 2015 *J. Chem. Theory Comput.* **11** 2087
- [159] Welborn M, Cheng L and Miller T F 2018 *J. Chem. Theory Comput.* **14** 4772
- [160] Cheng L, Welborn M, Christensen A S and Miller T F 2019 *J. Chem. Phys.* **150** 131103
- [161] Honig B and Karplus M 1971 *Nature* **229** 558
- [162] Warshel A and Levitt M 1976 *J. Mol. Biol.* **103** 227
- [163] Åqvist J and Warshel A 1993 *Chem. Rev.* **93** 2523
- [164] Mulholland A J, Lyne P D and Karplus M 2000 *J. Am. Chem. Soc.* **122** 534
- [165] Senn H M and Thiel W 2007 *Curr. Opin. Chem. Biol.* **11** 182
- [166] Magalhães R P, Fernandes H S and Sousa S F 2020 *Isr. J. Chem.* **60** 655
- [167] Zhang Y-J, Khorshidi A, Kastlunger G and Peterson A A 2018 *J. Chem. Phys.* **148** 241740
- [168] Shen L and Yang W 2018 *J. Chem. Theory Comput.* **14** 1442
- [169] Peterson A A 2016 *J. Chem. Phys.* **145** 074106
- [170] Jacobsen T, Jørgensen M and Hammer B 2018 *Phys. Rev. Lett.* **120** 026102