ARTICLE OPEN

Efficient training of ANN potentials by including atomic forces via Taylor expansion and application to water and a transition-metal oxide

April M. Cooper^{1,2}, Johannes Kästner², Alexander Urban^{1,3} and Nongnuch Artrith^{1,3}

Artificial neural network (ANN) potentials enable the efficient large-scale atomistic modeling of complex materials with near first-principles accuracy. For molecular dynamics simulations, accurate energies and interatomic forces are a prerequisite, but training ANN potentials simultaneously on energies and forces from electronic structure calculations is computationally demanding. Here, we introduce an efficient alternative method for the training of ANN potentials on energy and force information, based on an extrapolation of the total energy via a Taylor expansion. By translating the force information to approximate energies, the quadratic scaling with the number of atoms exhibited by conventional force-training methods can be avoided, which enables the training on reference datasets containing complex atomic structures. We demonstrate for different materials systems, clusters of water molecules, bulk liquid water, and a lithium transition-metal oxide that the proposed force-training approach provides substantial improvements over schemes that train on energies only. Including force information for training reduces the size of the reference datasets required for ANN potential construction, increases the transferability of the potential, and generally improves the force prediction accuracy. For a set of water clusters, the Taylor-expansion approach achieves around 50% of the force error improvement compared to the explicit training on all force components, at a much smaller computational cost. The alternative force-training approach thus simplifies the construction of general ANN potentials for the prediction of accurate energies and interatomic forces for diverse types of materials, as demonstrated here for water and a transition-metal oxide.

npj Computational Materials (2020)6:54; <https://doi.org/10.1038/s41524-020-0323-8>

INTRODUCTION

During the past decade, machine-learning potentials (MLP) trained on accurate first principles and quantum chemistry methods have become an integral part of computational materials science^{1–11}. Carefully constructed MLPs can be as accurate as their first principles reference method but at a fraction of the computational cost and with an effort that scales linearly with the number of atoms, which enables the modeling of complex materials that are not accessible with first principles methods such as density-functional theory (DFT)^{12–15}. MLPs can also be trained on results from highly accurate ab initio methods such as coupled-cluster calculations^{16–18}.

One popular MLP variant is the high dimensional artificial neural network (ANN) potential method initially introduced by Behler and Parrinello for elemental Si in 2007¹ and extended to multiple chemical species by Artrith and coworkers^{19,20}, exploring also the effect of long-ranged electrostatic interactions¹⁹. ANN potentials have been successfully used to model many complex materials, such as elemental metals^{21,22}, alloys^{23,24}, oxides²⁵, molecular systems^{26–29}, amorphous phases^{30–33}, interfaces^{34–37}, and nano-porous materials³⁸.

Once trained, the computational cost of ANN potentials does not scale with the number of data points used for training, so that training sets can be as large as necessary to sample the relevant and potentially diverse chemical and structural space. Often, ANN potentials are trained on total energy information, i.e., a single piece of information per DFT calculation is used as reference for the potential training. As a result, a large number of DFT reference

calculations, possibly tens to hundreds of thousands³³, may be needed to achieve the desired interpolation accuracy for applications in Monte Carlo (MC) sampling or molecular dynamics (MD) simulations.

Large training sets are needed to accurately capture the gradient of the potential energy surface (PES) and thus the interatomic forces. An accurate representation of the atomic forces is crucial for MD simulations and geometry optimizations, and hence the force prediction error is an important target parameter for potential construction. At the same time, the energy is what determines the most stable structure or phase and other materials properties, and energy conservation needs to be obeyed by any interatomic potential, so that an accurate representation of the structural energy is also needed.

In ANN potential training, the force prediction accuracy is typically converged by increasing the number of reference structures used for training until the relevant structure space is sampled sufficiently fine to also represent the gradient of the PES^{21,34}. This strategy does not only increase the training set size but also makes the training technically more challenging, since relevant structures have to be carefully selected without adding redundancies to the training set.

In principle, atomic forces or higher derivatives of the energy from first principles can also be used as reference data for ANN potential training. To train on force information, the loss function for the ANN potential training has to include the force prediction error, i.e., the error of the negative gradient of the energy. However, including the gradient in the loss function introduces a significant computational overhead because training requires the

¹School of Chemistry, University of St Andrews, St Andrews, UK. ²Institute for Theoretical Chemistry, University of Stuttgart, Stuttgart, Germany. ³Department of Chemical Engineering, Columbia University, New York, NY, USA. ✉email: nartrith@atomistic.net

Fig. 1 Flowchart of the construction of artificial neural network (ANN) potentials. **a** Often, ANN potentials are trained on total energies from reference electronic structure calculations, since direct training of gradients (interatomic forces) is computationally demanding. This approach requires large numbers of reference calculations to converge the slope of the potential energy surface. **b** In this work, a new scheme for ANN potential training is introduced, in which interatomic forces are used to extend the training set with energies approximated by Taylor expansion. The method is computationally as efficient as the conventional energy training and can significantly reduce the number of required reference calculations.

gradient of the loss function and thus the second derivative of the ANN potential is needed (i.e., the Hessian matrix). Therefore, in practice sometimes hybrid approaches are used in which the atomic forces of a subset of atoms are included in the loss function³⁹. Such approaches are especially useful in combination with online training methods that allow the selection of different force components for each training iteration (epoch).

In the present article, we introduce a new scheme for including atomic force information in the ANN potential training that avoids the computationally demanding evaluation of higher order derivatives (see flowchart in Fig. 1). The approach, which is based on a Taylor extrapolation of the total energy, is detailed in the following. In the Results section, we first demonstrate the basic principle of the methodology for an analytical example and then apply it to systems with increasing complexity (clusters of water molecules, bulk water, and a quaternary metal oxide), finding that a smaller number of reference structures compared to energy-only training is sufficient for converging the force error in ANN potential construction.

RESULTS

ANN potentials

ANN potentials are a type of many-body interatomic potential for atomistic simulations^{1,5,40}. In contrast to conventional potentials that are based on an approximate representation of the physical atomic interactions, such as embedded atom models⁴¹, ANN potentials employ general flexible functions, ANNs, for the interpolation between reference data points from first-principles calculations. ANN potentials represent the total energy $E(\sigma)$ of an atomic structure $\sigma = \{\vec{R}_i\}$ ($\vec{R}_i$ is the position vector of atom i) as the sum of atomic energies

$$E(\sigma) \approx E^{\text{ANN}}(\sigma) = \sum_i E_{\text{atom}}(\sigma_i^{R_c}) \quad (1)$$

where $\sigma_i^{R_c}$ is the local atomic environment of atom i , i.e., the atomic positions and chemical species of all atoms within a given cutoff radius R_c of atom i ^{1,5,40}. The atomic energy function E_{atom} in Eq. (1) is given by an ANN specific for each chemical species t (i.e., t is the type of atom i)

$$E_{\text{atom}}(\sigma_i^{R_c}) = \text{ANN}_t(\tilde{\sigma}_i^{R_c}) \quad \text{for atom type } t. \quad (2)$$

ANNs require an input of constant size, but the number of atoms within the local atomic environment of an atom i , $\sigma_i^{R_c}$, can vary. A suitable input feature vector is obtained by transforming $\sigma_i^{R_c}$ to a descriptor $\tilde{\sigma}_i^{R_c}$ with constant dimension that is also invariant with respect to (i) translation/rotation of the entire

structure, and (ii) exchange of equivalent atoms. In the present work, we employed the Chebyshev descriptors by Artrith, Urban, and Ceder²⁰, and the symmetry function descriptor by Behler and Parrinello^{1,42}.

Further details of the ANN architecture and the descriptor parameters are given in the Methods section.

ANN potential training with reference energies

ANN potentials are trained to reproduce the structural energies of reference datasets containing atomic structures σ (input) and energies $E(\sigma)$ (output) from first principles calculations. The atomic energy E_{atom} is not uniquely defined from first principles, so that no reference for the direct training of $\text{ANN}_t(\tilde{\sigma}_i^{R_c})$ is available. Note that there are non-unique approaches that can be used to decompose the total structural energy from first principles into atomic contributions⁴³. To avoid ambiguities and additional model complexity, it is more straightforward to implement the ANN potential training such that it minimizes a loss function based on the total energy $E(\sigma)$

$$\mathcal{L} = \sum_{\sigma} \frac{1}{2} [\Delta E(\sigma)]^2 \quad \text{with} \quad \Delta E(\sigma) = E^{\text{ANN}}(\sigma) - E(\sigma), \quad (3)$$

where $E^{\text{ANN}}(\sigma)$ is the energy of structure σ predicted by the ANN potential of Eq. (1).

Gradient-based optimization methods require the derivative of the loss function with respect to the ANN parameters, the ANN weights $\{w_k\}$

$$\frac{\partial}{\partial w_k} \mathcal{L} = \sum_{\sigma} \Delta E(\sigma) \frac{\partial}{\partial w_k} E^{\text{ANN}}(\sigma) = \sum_{\sigma} \Delta E(\sigma) \sum_{i \in \sigma} \frac{\partial}{\partial w_k} \text{ANN}_t(\tilde{\sigma}_i^{R_c}) \quad (4)$$

where the weight derivatives of the ANNs can be obtained using the standard backpropagation method⁴⁴.

The computational complexity of backpropagation scales as $\mathcal{O}(N_w)$ where N_w is the number of weight parameters. For a training set containing a total of N_{atom} atoms, the computational cost of one training epoch is therefore proportional to $\mathcal{O}(N_{\text{atom}} N_w)$.

Training with reference atomic forces

Density-functional theory (DFT)^{12,13} calculations with local or semi-local density functionals provide the interatomic forces at minimal overhead. Considering the importance of accurate forces for structure optimizations and MD simulations, it is desirable to include the force error in the loss function $\mathcal{L}$ of Eq. (3). In addition, each atomic structure has only one total energy but three atomic force components per atom. Hence, using the atomic forces as

additional reference data increases the data points available for training per atomic structure by a factor of $3N$, where N is the number of atoms in the structure.

The Cartesian vector of the force acting on atom i is given by the negative gradient of the total energy

$$\vec{F}_i(\sigma) = -\vec{\nabla}_i E(\sigma) \quad \text{with} \quad \vec{\nabla}_i = \frac{\partial}{\partial \vec{R}_i} \quad (5)$$

where $\vec{R}_i = (R_i^x, R_i^y, R_i^z)^T$ is the Cartesian position vector of atom i . Including the atomic forces in training can thus be accomplished with the loss function

$$\begin{aligned} \mathcal{L}' &= \mathcal{L} + a \sum_{\sigma} \frac{1}{2} \left[\sum_j \left(-\vec{\nabla}_j E^{\text{ANN}}(\sigma) - \vec{F}_j(\sigma) \right) \right]^2 \\ &= \mathcal{L} + a \sum_{\sigma} \frac{1}{2} \left[\Delta \vec{F}(\sigma) \right]^2 \quad \text{with} \quad \Delta \vec{F}(\sigma) = \sum_j \left(-\vec{\nabla}_j E^{\text{ANN}}(\sigma) - \vec{F}_j(\sigma) \right), \end{aligned} \quad (6)$$

where a is an adjustable parameter that determines the contribution of the force error to the loss function $\mathcal{L}'$. The gradient of the new loss function $\mathcal{L}'$ with respect to the weight parameters is

$$\frac{\partial}{\partial w_k} \mathcal{L}' = \frac{\partial}{\partial w_k} \mathcal{L} - a \sum_{\sigma} \left[\sum_{j \in \sigma} \left(-\vec{\nabla}_j E^{\text{ANN}}(\sigma) - \vec{F}_j \right) \right] \sum_{j \in \sigma} \frac{\partial}{\partial w_k} \vec{\nabla}_j E^{\text{ANN}}(\sigma). \quad (8)$$

As seen in Eq. (8), evaluating the weight gradient of the new loss function $\mathcal{L}'$ requires taking the derivative of the position gradient of the ANN potential

$$\sum_{j \in \sigma} \frac{\partial}{\partial w_k} \vec{\nabla}_j E^{\text{ANN}}(\sigma) = \sum_{j \in \sigma} \sum_{t \in \sigma} \frac{\partial}{\partial w_k} \vec{\nabla}_j \text{ANN}_t(\vec{\sigma}_i^{R_c}) \quad (9)$$

which scales quadratic with the number of atoms in the reference structure. Note that only the cross terms for atoms within two times the cutoff radius of the potential ($2R_c$) are different from zero²¹. For very large or periodic structures the scaling of the derivative evaluation therefore eventually becomes linear but with a large prefactor of N_{local} that is the average number of atoms within $2R_c$ and can be several hundred to thousands depending on the density of the material and the cutoff radius.

The quadratic scaling and the often large number of atoms in reference datasets typically makes it infeasible to train all force components using the loss function in Eq. (7). One option is to include a fraction of all force components in the loss function, e.g., only 0.41% of the atomic forces were included in ref. ³⁹, and Artrith et al. have previously used only 10% or less of the force information for ANN potential training^{19,21,34}. This strategy can work especially well with online training methods that allow selecting different force components at each epoch, which is not possible using batch training methods.

The importance of atomic force information and the unfavorable scaling of direct force training prompted us to consider alternative means of including force information in ANN potential training.

Including atomic force information via Taylor expansion

Training of the energy with the loss function of Eq. (3) is efficient, so that a translation of the force information to (approximate) energy information is advantageous. Such a translation can be accomplished using a first-order Taylor expansion to estimate the energy of additionally generated atomic structures without the need to perform additional electronic structure calculations.

The energy of a structure $\sigma' = \{\vec{R}_i'\}$ that was generated by displacing the atoms in the original structure $\sigma = \{\vec{R}_i\}$ can be

expressed as the Taylor series

$$E(\sigma') = E(\sigma) + \sum_i \vec{\delta}_i \vec{\nabla}_i E(\sigma) + \frac{1}{2} \sum_i \vec{\delta}_i^2 \vec{\nabla}_i^2 E(\sigma) + \dots \quad \text{with} \quad \vec{\delta}_i = \vec{R}_i' - \vec{R}_i. \quad (10)$$

Substituting the atomic force of Eq. (5) for the negative gradient of the energy and truncating after the first order, we arrive at the approximation

$$E(\sigma') \approx E(\sigma) - \sum_i \vec{\delta}_i \vec{F}_i(\sigma), \quad (11)$$

where $E(\sigma)$ and $\vec{F}_i(\sigma)$ are the energies and atomic forces from the original reference electronic structure calculations. The first-order approximation is valid for small displacements $\{\vec{\delta}_i\}$.

Equation (11) provides a mechanism for incorporating approximate force information in the ANN potential training as additional structure-energy pair records. What remains is to decide on a recipe for generating such additional structures by atom displacement. In the present work, we considered two strategies: (A) displacing one individual atom to generate one additional structure, and (B) randomly displacing all atoms. In strategy (A), we select an atom i in structure σ and displace its coordinates by a small amount $\pm \delta$ in each Cartesian direction to generate six additional structures. For example, displacing atom i in structure $\sigma = \vec{R}_1, \vec{R}_2, \dots, \vec{R}_N$ in the negative Cartesian x direction would yield the new structure-energy pair

$$\begin{aligned} \sigma' &= \{\vec{R}_1, \dots, \vec{R}_i - \delta \hat{x}, \dots, \vec{R}_N\} \\ E(\sigma') &= E(\sigma) - \delta F_i^x(\sigma), \end{aligned} \quad (12)$$

where $\hat{x}$ is the unit vector in the Cartesian x direction, and F_i^x is the x component of the force acting on atom i . A similar approach has previously been used by Vlcek et al. for molecular force-field optimization with energy and force information⁴⁶.

In strategy (B), all atoms are displaced by small random vectors $\vec{\delta}_i$ where the total displacement is such that $|\vec{\delta}_i| \leq \delta_{\text{max}}$ for all atoms. Since a net translation of the entire structure does not affect the energy, the center-of-mass displacement is subtracted from the combined atomic displacements. The energy of the resulting structure is evaluated according to Eq. (11).

The optimal values of the displacement δ in strategy (A) and the maximal displacement δ_{max} in strategy (B) are parameters and will be determined in the following.

To assess the efficacy of the Taylor-expansion approach laid out in the previous section, we considered a series of materials systems with increasing complexity: An analytic Lennard-Jones dimer molecule as test case, clusters of water molecules, a periodic bulk water box, and a complex oxide system with five different chemical species.

Diatomic molecule

Our first test case is a diatomic molecule with atomic interactions described by the analytic Lennard-Jones potential⁴⁷

$$V(r) = \varepsilon \left[\left(\frac{r_0}{r} \right)^{12} - 2 \left(\frac{r_0}{r} \right)^6 \right] \quad (13)$$

with binding energy $\varepsilon = 3.607$ eV and equilibrium distance $r_0 = 1.54$ Å. The binding energy and bond distance were chosen such that they correspond to a typical covalent bond, approximately the carbon-carbon bond, so that the magnitude of the displacement in the Taylor expansion will be comparable to an actual compound. An analytic interatomic potential has the advantage that the gradient and the Taylor expansion can be calculated analytically. Further, the one-dimensional dimer molecule makes it straightforward to visualize the PES, i.e., the bond-energy curve, so that the analytic potential and the ANN interpolation can be visually compared with each other.

Fig. 2 Analytic dimer example illustrating the Taylor expansion approach. The gray lines indicate the energy vs. bond length curve of the model Lennard–Jones potential discussed in the text, and the seven samples used for ANN potential training are shown as blue circles. The red line in panel **a** corresponds to an ANN potential trained only on the energy of the reference data points (blue circles). In panel **b** the ANN potential (red line) was additionally trained on 14 displaced structures (black crosses) with displacements of $\delta = \pm 0.02$ Å, the energies of which were approximated by a first-order Taylor expansion using the analytic gradient at the blue reference data points.

Figure 2a shows the bond energy in a range close to the equilibrium distance and an ANN potential that was trained on seven equidistant reference points between 1.4 Å and 2.8 Å. As seen in the figure, the ANN potential approximates the Lennard–Jones potential well at distances above 1.8 Å, but is unable to reproduce the minimum region and predicts an incorrect larger equilibrium distance as well as an incorrect slope and curvature near the minimum. The result of training the ANN potential with the Taylor expansion approach is shown in Fig. 2b, where additional data points were generated by approximating the energy for slightly longer and slightly shorter bond distances using the first-order Taylor expansion of Eq. (11) using a displacement of $\delta = 0.02$ Å. As seen in the figure, the additional approximate data points guide the ANN potential, and the resulting fit is visually much better and reproduces the correct bond minimum as well as the slope and the curvature within the minimum well.

Also apparent in Fig. 2b are the deviations of the first-order Taylor expansion from the true bond energy curve. The approach is approximate and will give rise to noise in the reference data, and the optimal amount of additional approximate data points per exact energy data point has to be determined such that the accuracy of the ANN potential does not suffer. In the following, we will assess the Taylor expansion method for real materials systems to quantify both the improvement of the force prediction accuracy and the effect of noise in the reference energies.

Water clusters

While the analytic dimer example is useful for illustrative purposes, our objective is the training on first-principles energies and forces. As a second test case, we therefore consider a reference dataset based on structures obtained from MD simulations of water clusters with six water molecules. MD simulations at 300 K and 800 K were performed on the semiempirical Geometry, Frequency, Noncovalent, eXtended TB (GFN-xTB)⁴⁸ level of theory. The atomic forces and energies of a subset of the structures along the MD trajectories were recalculated using a first-principles density-functional theory (DFT) approach (BLYP-D3/def2-TZVP) to be used as reference data for ANN potential construction. See the Methods section for the details of our MD and DFT calculations.

As a baseline for the assessment of the force-training method, we first quantify the error in the atomic forces if only total energies are trained using the methodology and loss function for energy training described at the beginning of the Results section. As a second point of reference the error in the atomic forces is

quantified for ANN potentials that were trained by including the error in the atomic forces in the loss function as described above.

Total energy training with increasingly large training sets. As mentioned in the introduction, a common way to improve the force prediction error of ANN potentials is by increasing the reference dataset size to sample the configurational space more finely. We therefore investigated first the influence of the size of the reference dataset on the quality of the force prediction of ANN potentials that were trained on the total energy only.

In order to study how the size of the reference dataset affects the quality of the force prediction, three different reference datasets with increasing number of data points were assembled from the MD reference dataset:

- (i) a subset containing 471 reference structures referred to as the `train_0500` dataset in the following,
- (ii) a set with 943 structures (`train_1000`), and
- (iii) a set with 1886 structures (`train_2000`).

The structures within these subsets were chosen evenly spaced along the MD trajectories to ensure a maximum decorrelation of the reference data.

For each training set, the ANN potential training was repeated 10 times with different random initial weight parameters $\{w_k\}$ to obtain statistics on the prediction quality of the atomic forces for the resulting ANN potentials. Details on the ANN potential training are given in the methods section. All errors reported in the following were obtained for the same independent validation dataset containing 2000 structures that are not included in any of the training sets.

Figure 3a shows the distributions of the error in the norm of the predicted atomic forces for the structures within the validation set after training on the energies in the `train_0500`, `train_1000`, and `train_2000` datasets. Only the energy error entered the loss function, so that the interatomic forces were not directly trained.

As seen in the figure, for the smallest `train_0500` set interpolating the ANN PES using only the energies of the reference structures leads to a wide distribution of errors in the prediction of the absolute value of the atomic forces. Especially, the pronounced tail of the distribution implies that there is a large fraction of atoms for which the absolute value of the force is predicted with an error greater than 1.0 eV/Å. Increasing the size of the training set to ~1000 and ~2000 structures reduces the tail of the error distribution significantly.

Figure 3b–d shows a corresponding analysis of the error in the direction of the predicted atomic forces and the mean absolute

Fig. 3 Impact of the training set size on the ANN force prediction errors for clusters of six water molecules. **a** Force error distribution after training on the three different training sets with ~500 (*train_0500*, red shaded), ~1000 (*train_1000*, dark red stars), and ~2000 (*train_2000*, violet squares) structures. The inset shows a representative water cluster structure. **b–d** Frequency of occurrence of a given error in the direction of the force in degrees as a function of the absolute atomic force for energy training and varying training set size. Panels **b–d** also show the mean absolute errors (MEA) of the predicted forces for the three different training sets. All statistics shown are based on 10 ANN potentials trained for each scenario.

error (MAE) of the atomic forces. High relative frequencies of occurrence are shown in yellow and red, whereas low relative frequencies are colored in shades of gray. The errors are shown as a function of the absolute value of the atomic force, and it can be seen that the reliability of the prediction of the direction of the atomic forces increases for increasing absolute values of the force vectors. This means, the error in the force direction is greater for small force vectors than for large force vectors. Especially for atoms with atomic forces with absolute values of less than 1.0 eV/Å, the direction of the force vector predicted by the ANN potential scatters strongly. This scattering is significantly decreased for atomic forces with large absolute values. As seen in panel a, for the small *train_0500* set the error distribution has a shallow maximum between 0 and 25° depending on the absolute force value, but the heat map shows much larger errors of nearly 180° for force vectors with small absolute value.

Additionally, increasing the size of the reference dataset reduces the scattering in the predictions notably. Particularly, the scattering in the prediction of atomic forces with absolute values smaller than 1.0 eV is notably decreased in comparison to the results obtained from the *train_0500* reference dataset. Furthermore, the number of atoms for which the error is 15° or less increases significantly when the ANN potential is trained with the *train_2000* reference data instead of the *train_0500* dataset.

Depending on the reference method and the size of the atomic structures, increasing the number of structures in the training set may entail a massive computational overhead. Therefore, we next investigate whether the Taylor-expansion formalism could alleviate the need for large training set sizes.

Optimal meta-parameters for the Taylor-expansion approach. Next, we investigate the efficacy of approximate force training using the Taylor-expansion approach. In order to apply the two displacement strategies, i.e., (A) the displacement of single atoms in the three Cartesian directions and (B) the displacement of all atoms in random directions, suitable displacement parameters δ and $\delta_{\max}$ have to be determined. Finally, the optimal number of additional structures with approximate energies to be generated using the Taylor-expansion approach needs to be determined, as the computational effort for ANN potential training scales with the size of the training set.

The number of additional structures is given in terms of a multiple a of the original training set, so that $a = X$ means that the number of generated structures is X times the number of the original structures. For example, the *train_0500* dataset contains 471 structures, so that for $a = 10$ a total of $10 \times 471 = 4710$

Fig. 4 Relative mean absolute error (MAE) of the atomic forces for different force-training parameters. The MAE is shown as a function of the atomic displacement $\delta_{\max}$ for different multiples a of additionally generated structures. The graph shows the MAE relative to the DFT reference. Results are shown for displacement strategy (B) and for the smallest reference dataset (*train_0500*). The red dashed line indicates the reduction of the force error that can be achieved by direct force training.

additional structures with approximate energy will be generated by atomic displacement.

For each possible choice of a parameter pair (a, δ) or $(a, \delta_{\max})$ 10 ANN potentials were trained to obtain statistics on the resulting ANN potentials. The optimal parameter pair was chosen by calculating the MAE of the atomic forces for the validation dataset and thereby averaging the errors obtained from all 10 potentials fitted for each parameter pair.

Figure 4 shows the MAE of the atomic forces obtained after training with the Taylor-expansion approach as function of the maximum displacement $\delta_{\max}$ for different multiples of additional structures a . The ANN potentials used to quantify this error were fitted to the *train_0500* reference dataset using random displacement strategy (B) for the given $\delta_{\max}$ parameters. In the figure, the MAE are given relative to the one obtained if no force information is used for training.

As seen in Fig. 4, the values of $\delta_{\max}$ that lead to the greatest reduction of the MAE are independent of the number of additionally generated structures a and are close to $\delta_{\max} = 0.01$ Å. If the displacement parameter is chosen too large, the first-order Taylor expansion is no longer a good approximation and the force error increases. Additionally, the MAE decreases with increasing number of generated structures until a multiple of $a = 22$ has been reached. For multiples a greater than 22 no

Fig. 5 Impact of force training on the ANN force prediction errors for clusters of six water molecules. **a** Force error distribution for training on the `train_0500` training set only with energy information (Energy training, red shade) and using approximate force training with the Taylor-expansion approach and the Cartesian (green crosses) and random (blue triangles) displacement strategies. **b** The same data as in panel **a** but in addition the error distribution resulting from direct (exact) force training is shown as thick black line. **c–e** Frequency of occurrence of a given error in the direction of the force in degrees as a function of the absolute atomic force, using approximate force training with the Taylor-expansion approach. The optimal values for the displacement and the multiple a were used (Cartesian displacement: $\delta = 0.03$ Å, $a = 11$, Random displacement: $\delta_{\max} = 0.008$ Å, $a = 22$). All statistics shown are based on 10 ANN potentials trained for each scenario.

significant further improvement was found. Based on this analysis we conclude that, for the `train_0500` reference dataset, the optimal parameter choice that leads to the smallest MAE with the Taylor-expansion approach is $a = 22$ and $\delta_{\max} = 0.008$ Å.

The optimization of the meta-parameters was repeated for the larger training sets `train_1000` and `train_2000` and for the second Taylor-expansion displacement strategy (A), and the optimal parameters are given in Table S1 in the Supplementary Information. In summary, the optimal displacement δ and multiple a do not appear to depend on the size of the reference dataset. Out of the considered values, the optimal multiple of generated structures is $a = 22$, and optimal displacements are $\delta_{(A)} = 0.03$ Å for strategy (A) and $\delta_{(B)} = 0.008$ Å for strategy (B). The heatmaps in Supplementary Fig. 1 show the similarity of the error in the force direction for varying displacements.

Energy and force training with the Taylor-expansion approach. Using the optimal parameters for the displacement δ and fraction of generated structures a for each of the three reference datasets and both displacement strategies, we studied the errors in the prediction of the atomic forces in order to compare both displacement strategies with each other, and to quantify the improvement of the predicted forces with respect to the conventional approach where only total energies of the reference structures are fitted.

Figure 5a compares the distribution of absolute force errors in the validation set after training on the `train_0500` set with and without force information. As seen in the figure, the errors in the absolute force are drastically reduced by both displacement strategies. Especially the tail of the error distribution with errors

greater than 1.0 eV/Å has nearly disappeared for the ANN potentials that were trained using the Taylor-expansion approach. In addition, the prediction of the direction of the force vectors also improves distinctly, as seen in Fig. 5c–e. Comparing the results obtained with the Taylor-expansion approach to those for conventional energy training (Fig. 3b–d), a general decrease of the error in the predicted forces can be observed that is also reflected by a reduction of the MAE.

Based on the error distributions in Fig. 5a, both displacement strategies perform nearly equally well. Displacing atoms along the Cartesian directions (strategy A) generating 5088 additional structures results in a slightly smaller improvement than displacing atoms in random directions (strategy B) generating 10,362 structures, but both Taylor-expansion strategies are a significant improvement over energy-only training. For the small water cluster system, the Cartesian displacement strategy requires generating far fewer additional structures than the random displacement approach. We found that this is not generally the case and instead depends on the number of atoms in the reference structures. With increasing structure size, the optimal multiple a increases more rapidly for displacement strategy (A), and thus for larger structures displacement strategy (B) becomes more efficient.

Comparison of Fig. 5a with Fig. 3a and the heatmap in Fig. 5c with that in Fig. 3d shows that training the `train_0500` dataset with approximate force information results in quantitatively similar error distributions as training on the energies in the four-times larger `train_2000` dataset. Hence, the approximate force training reduces the number of training data points by a factor of four for the example of the water cluster dataset.

In principle, both the size of the reference dataset as well as the force training can be expected to affect also the error in the predicted energy. In the case of the water cluster dataset, even training on the energies of the `train_0500` dataset without force information results in energy errors in the validation set that are well below 2×10^{-3} eV/atom, which is an order of magnitude below chemical accuracy (1 kcal/mol $\approx$ 0.04 eV). Therefore, the water cluster dataset is not well suited to investigate the impact of force training on the energy prediction, and we will return to this question when discussing more complex datasets in the following section. For the water cluster dataset, Supplementary Fig. 2 shows that the distribution of the energy error is not significantly affected by force training using the Taylor-expansion approach.

Comparison of approximate and direct force training. For the small water cluster structure (18 atoms) and the smallest (`train_0500`) dataset, direct force training by including all force errors in the loss function is computationally feasible. Note that even for this simple system, the direct force training took on average eight times more computer time per iteration (training epoch) than the Taylor-expansion method with data multiple $a = 22$ on our computer system. The force error distribution resulting from direct force training is shown in Fig. 5b.

Comparing the results obtained by approximate force training to those from direct force training, it is obvious that direct force training improves the force prediction even further. The tail of the distribution of the absolute force error approaches zero at around 0.5 eV/Å, and as a result the error distribution is significantly narrower than the ones obtained by the other methods with its maximum being closer to 0.

The same trend is also found for the improvement of the prediction of the direction of the force vector by direct force training compared to the approximate Taylor-expansion approach. The mean absolute force error for direct force training is 0.18 eV/Å, and a heatmap of the direction error is shown in Supplementary Fig. 3.

Quantitatively measured by the MAE, approximate force training using the Taylor-expansion methodology gives around half the reduction of the force error that is achieved by direct force training. This can be most clearly seen in Fig. 4, in which the relative MAE obtained after direct force training is indicated by a red dashed line.

Even though direct force training leads to yet better predictions of the forces, one has to take into account that it is significantly more computationally demanding. For larger and more complex systems for which direct force training would be challenging or even infeasible, approximate force training with the Taylor-expansion approach offers an alternative to improve the force prediction significantly in comparison to conventional energy training. Therefore, we investigate condensed phases in the following sections.

Bulk water

We applied the Taylor-expansion methodology to a reference dataset of a periodic bulk water system with 64 water molecules (192 atoms) that was generated by running an ab initio MD (AIMD) simulation at a temperature of 400 K over a simulation time of 100 ps with time steps of 1 fs, producing a total of 100,000 MD frames. The simulations employed a Γ -point k-point mesh for the numerical Brillouin-zone integration. All parameters of the DFT calculations are detailed in the methods section.

Seven hundred structures along the AIMD trajectory were selected for potential training. Ninety percent of the reference dataset were used for ANN potential training, and the remaining 10% were used as an independent test set to monitor training progress, and to detect overfitting. Two thousand different equally spaced MD frames from the complete trajectory were compiled into a third independent set for validation, and all results within

this section are based on the validation set none of which was used for training.

Since the nature of the bonds in the water bulk is the same as in the water clusters of the previous section, a maximal displacement of $\delta_{\max} = 0.01$ Å was used for the water bulk system, which was found to be close to optimal for water clusters (see Fig. 4). Additionally, we limit the discussion to random displacement strategy (B), since it performed better than displacing individual atoms for water clusters. Instead, here we focus on the impact of increasing the number of additional structures generated by the Taylor-expansion approach.

In general, two competing effects can be expected: On the one hand, the force training should become more effective with increasing multiple a , i.e., with an increasing number of additional structures per original structure generated by our approach. On the other hand, the first-order Taylor-expansion is approximate, and the energies of the additional structures are less accurate and noise is introduced into the training set. Thus, the accuracy of the energy fit could be expected to decrease after introducing too many additional structures via atomic displacement.

To determine the balance of these two effects, data multiples between $a = 12$ and $a = 64$ were considered, which corresponds to $\sim 12 \times 700 = 8400$ to $64 \times 700 = 44,800$ additional structures generated by Taylor expansion for the 700-structure reference dataset. See the Methods section for details of the ANN potential construction.

Figure 6a, b shows, as solid lines, the best energy and force prediction errors out of 10 ANN potentials trained on the same data but with different initial weight parameters. Additionally, the median ANN potential error is shown (dashed lines) as a proxy for the likely result of a single ANN potential training run. As seen in the figures, both the energy and force errors initially decrease with increasing amount of additional data. For small data multiples a , the force error improves significantly from 1.26 eV/Å (energy training only) to 0.88 eV/Å (30% improvement) for $a = 12$ and 0.69 eV/Å (45%) for $a = 24$ (Fig. 6b). Increasing a beyond 24 only yields marginal further improvement, and for $a = 48$ the force error has decreased to 0.61 eV/Å (52%).

Interestingly, the energy RMSE decreases simultaneously from 2.6 meV/atom (energy training) to 2.1 meV/atom (19%) for $a = 36$ (Fig. 6a). This improvement indicates that the displaced structures added to the training set via Taylor expansion have improved the transferability of the resulting ANN potential. Increasing a further does not result in further improvement of the energy error, and an increase of the error is observed for $a = 64$, which is in-line with our expectation that too many additional structures introduce noise in the training set.

Having quantified the general improvement of the predicted atomic forces with the force-training approach, we analyze next where these improvements originate from. Figure 7a shows the distribution of absolute force errors in the validation set with and without force training. In the case of energy training only, the force errors are widely spread out, and for some atoms the force error is greater than 3.0 eV/Å. The largest atomic forces in the dataset are around 5.0 eV/Å, so that 3.0 eV/Å corresponds to an error of at least 60%. As discussed above, the force prediction error depends on the size of the reference dataset, and thus the analysis in Fig. 7a shows that the 700-structure dataset is not sufficient for robust force prediction.

Force training with the Taylor-expansion approach results in a strong improvement of the force prediction, and training with a data multiple of $a = 16$ already removes the high-error tail of the error distribution. Increasing a to 36 further reduces the width of the distribution significantly. The error distributions for a large data multiple of $a = 64$ is nearly identical to the distribution for $a = 36$ and both are centered around 0.6 eV/Å, showing that the force training has converged for this dataset in agreement with Fig. 6.

Fig. 6 Change of the energy and force error with increasing number of additional structures from Taylor expansion for bulk water. **a** Root mean squared error (RMSE) of the energy relative to the DFT reference energies for ANN potentials trained with and without force information. The solid line shows the error of the best potential out of 10 training runs, and the dashed line indicates the corresponding median. **b** The equivalent analysis for the error in the norm of the interatomic forces. A representative bulk water structure is shown as inset in panel **a** (oxygen atoms are red, and hydrogen is white).

Fig. 7 Error in the force norm and direction for bulk water. **a** Distribution of the absolute error of the atomic forces without force training (red area) and with force training using increasingly more additional structures generated by displacement. **b** Distribution of the error in the force direction without (left) and with (right) force training. For the force training, a maximal displacement of $\delta_{\max} = 0.01$ Å and a structure multiple of $a = 64$ were used.

Figure 7b shows a heatmap of the distribution of the errors in the direction of the atomic forces for all atoms in the validation set for training without forces (left) and with forces (right). The results for $a = 64$ are shown in the figure, though $a = 36$ yielded a qualitatively equivalent error distribution. As seen in the figure, when training only energies the errors in the force direction are spread out over nearly the entire angle range. In contrast, the force training improves the error distribution significantly, yielding a maximum at around 15° and a strong decay of the error with the norm of the force vectors.

As another test, we constructed preliminary ANN potentials trained with and without force information on a larger reference dataset of 10,000 frames taken from the same AIMD trajectory to evaluate the radial distribution function (RDF) of liquid water. In Fig. 8, the results are compared to a literature reference for an equivalent DFT approach⁴⁹. Note that the construction of robust ANN potentials for water requires datasets with very diverse structures and phases that include also short O–H bond lengths that occur rarely in MD simulations^{26,29,50}, and the preliminary potentials can only be considered a starting point for the ANN potential construction and will require further refinement.

As seen in Fig. 8, the ANN potential trained with the Taylor-expansion approach yielded improvements in the RDF peak of the first coordination shell. Importantly, the stability of the MD simulations improved significantly with force training. Simulations using potentials trained on energies only were unstable due to frequent extrapolation, whereas the potentials trained with force information allowed for more robust MD simulations. It is clear

Fig. 8 O–O radial distribution functions (RDF) for bulk water. The RDF was calculated using preliminary ANN potentials trained with the Taylor-expansion approach (red line) and using energies only (blue line), respectively. The gray region indicates results from ab initio MD simulations with the same density functional at 300 K from ref.⁴⁹. The ANN-potential RDFs were evaluated for simulation cells with 128 water molecules.

that the potential that was trained using forces via the Taylor-expansion method performed better than the one that was trained on energies only.

Direct force training of all atomic force components was not feasible for the bulk water system. As detailed above, for

Fig. 9 Approximate force training applied to a lithium transition-metal oxide. Change of **a**, the root mean squared error (RMSE) of the energy and **b**, the mean error in the force norm as function of the maximal displacement $\delta_{\max}$. **c** Distribution of the absolute force error for different training parameters. **d** Distribution of the error in the force direction. A representative crystal structure of the Li–Mo–Ni–Ti oxide is shown as inset of panel **a**.

condensed phases direct force training scales with the number of atoms within twice the potential cutoff radius. For $R_c = 6.5$ Å, there are on average 960 atoms within range, so that direct force training of all atomic forces would take ~ 1000 times the computer time of energy-only training.

As shown above, the Taylor-expansion approach for force training works robustly for both isolated water clusters and water bulk. For materials with more complex compositions, sampling the structural space with sufficient resolution is challenging, and force training becomes even more important, as discussed in the next section.

Quaternary metal oxide

To determine the performance of the force-training approach for a material with complex chemical composition we finally investigated a quaternary transition metal oxide. Our benchmark system, the Li–Mo–Ni–Ti oxide (LMNTO), is of technological relevance as prospective high-capacity positive electrode material for lithium-ion batteries⁵¹. The compound exhibits substitutional disorder in which all four metal species, Li, Mo, Ni, and Ti, share the same sublattice.

We generated a reference dataset for LMNTO with composition $\text{Li}_8\text{Mo}_2\text{Ni}_7\text{Ti}_7\text{O}_{32}$ by running a 50 ps long AIMD simulation at 400 K as described in detail in the methods section, yielding a total of 50,000 MD frames. Again, training (~ 720 structures), test (~ 80 structures), and validation (1800 structures) sets were generated. A representative structure of the LMNTO unit cell is shown as inset in Fig. 9a.

The bonding in lithium transition-metal oxides exhibits mostly ionic character, and the bond strength cannot be expected to be the same as in the water systems of the previous sections. Therefore, the maximal displacement $\delta_{\max}$ for the Taylor-expansion approach first needs to be optimized for LMNTO.

Figure 9a, b shows the change of the energy and force errors as function of $\delta_{\max}$ for a fixed additional data multiple of $a = 70$. As seen in subfigure b, the mean error in the predicted force decreases with increasing maximal displacement and has not yet

converged for $\delta_{\max} = 0.04$ Å, which is four times greater than the optimal displacement found for water. On the other hand, the median energy RMSE has a minimum at around $\delta_{\max} = 0.02$ Å beyond which the energy error increases. A maximal displacement of $\delta_{\max} = 0.03$ Å is a good compromise for LMNTO, resulting in a small decrease in accuracy for the energy from 4.6 meV/atom to 4.9 meV/atom (increase by 6.5%) but a strong improvement in the accuracy of force prediction from 0.92 eV/Å to 0.66 eV/Å (28.3%).

The distribution of errors in the absolute values and directions of the atomic forces are shown in Fig. 9c, d, respectively. The results for $\delta_{\max} = 0.03$ Å are shown. Qualitatively, the same improvement of the absolute force error as in the case of the water systems is seen, though even without force training the error distribution does not exhibit any high-error tail. The force training also results in an improved accuracy in the direction of the predicted forces, though the maximum of the error distribution is at a slightly higher value of around $\sim 20^\circ$.

Towards application in MD simulations of LMNTO. The MAE of the forces and the RMSE of the energy are abstract quality measures of the ANN potential, but in practice the robustness and reliability of the potential in an actual application is most important. To construct a preliminary ANN potential for LMNTO, we compiled a dataset of ~ 4000 AIMD frames that were used as reference for training three ANN potentials with force information ($\delta_{\max} = 0.015$ Å, $a = 20$) and three without force information. The construction of accurate ANN potentials for materials with complex structures and compositions typically requires training sets with tens of thousands of reference structures^{24,33,34}. Four thousands reference structures taken from a single AIMD trajectory can function as an initial dataset for the construction of a preliminary ANN potential as a starting point for subsequent iterative refinement^{5,40}.

The robustness of an interatomic potential and the smoothness of the predicted PES is reflected by the numerical energy conservation in MD simulations. To test energy conservation, we carried out MD

Fig. 10 Energy conservation in microcanonical molecular dynamics (MD). **a** Representative LMNTO structure taken from MD simulation (colors as in Fig. 9). **b** Change of the total energy during MD simulations in the microcanonical (*NVE*) statistical ensemble using ANN potentials trained on energy information only (orange and red lines) and trained using the Taylor-expansion approach (green line).

simulations in the microcanonical (*NVE*) statistical ensemble using a time step of 1.0 fs, and the resulting change of the total energy over the course of 1.0 ns is shown in Fig. 10 for different ANN potentials. The MD simulations were performed for a $2 \times 2 \times 1$ LMNTO supercell containing 224 atoms (Fig. 10a) that was previously thermally equilibrated at a temperature of 400 K with AIMD simulation. All of the ANN potentials trained with approximate force information conserved the total energy well with fluctuations on the order of 10^{-3} meV/atom (one example is shown in Fig. 10b). However, those potentials that were trained only on the energies showed numerical instabilities of varying degree, and two examples are plotted in Fig. 10b.

This experiment demonstrates that the additional structure-energy data points generated from force information using the Taylor-expansion approach contribute to the transferability of the ANN potential and improve the smoothness of the potential interpolation.

DISCUSSION

We introduced a computationally efficient method for the simultaneous training of energies and interatomic forces for the construction of accurate ANN potentials. We assessed the methodology for three relevant complex materials systems: water cluster structures, bulk liquid water, and a solid metal oxide. We demonstrated that the approach increases the accuracy of predicted forces, reduces the number of first principles reference data points needed, and can improve the transferability of ANN potentials.

In general, we find that force training has the greatest impact for small reference datasets. With force training, a small reference dataset of ~500 water cluster structures could achieve nearly the same predictive power as training only the energy on a reference dataset with ~2000 structures. Using the force-training approach, the structural and chemical space can thus be more coarsely sampled than using energy-only training. This has important implications for the construction of ANN potentials for materials with complex compositions, such as the quaternary metal oxide of the previous section, for which an exhaustive sampling of the structural and chemical space is infeasible.

The Taylor-expansion methodology is approximate, and it is meant as a computationally more efficient alternative to the direct training of force information with a modified loss function. Direct force training also incurs a significant memory overhead if the derivatives of the atomic-structure descriptors are stored in memory, whereas the Taylor-expansion approach uses exactly the same amount of memory as training with energy information only. As demonstrated for the water cluster dataset, direct force training can further improve the force prediction when it is computationally feasible. However, even for the small water

cluster system, direct force training was already eight times more computationally demanding than approximate force training, and owing to its formal quadratic scaling this difference will be even greater for larger systems.

We discussed two different strategies for the generation of additional approximate reference energy data points from force information by approximating the energy of structures with slightly displaced atoms using a first-order Taylor expansion. The two strategies differ in the number of atoms that are displaced, but both rely on randomization, either for the selection of atoms or to decide the direction and magnitude of atomic displacements. Displacing all atoms in a reference structure by small random vectors with maximal length $\delta_{\max}$ is shown to be a robust scheme for the generation of derived structures. However, this strategy does, in principle, not rule out that multiple additional structures with similar information content (i.e., similar displacements) are generated. The generated atomic displacements are not necessarily linearly independent.

The atomic forces of a structure with N atoms comprise $3N$ pieces of information from the three force components of all atoms. Hence, at most $3N$ structures with independent information content can be derived from any given structure with energy and force information. One set of generated structures with linearly independent atomic displacements is given by the set of $3N$ normal modes^{52,53}. Excluding the three translations and three rotations that do not change the energy, one could employ displacements in the remaining $(3N - 6)$ normal mode directions as a third strategy for our force training approach. Normal-mode sampling has previously been shown to be a useful strategy for the generation of reference datasets, especially for molecular systems^{17,27}. However, we found empirically that far fewer than $3N$ additional structures are required to converge the error of the interatomic forces for a given reference dataset. For the bulk water system with 64 molecules, $N = 3 \times 64 = 192$, so that the number of normal modes is $3 \times 192 - 6 = 570$, but our computational experiments show that the force error already plateaus for 24–48 additional structures with random atomic displacements. An exhaustive enumeration of all degrees of freedom would therefore be suboptimal for our force-training methodology, and it is not obvious which normal modes should be selected if a subset was to be used. The random displacement strategy offers a reasonable compromise of information density and generality.

As evident from the water clusters and the oxide system, the optimal maximal displacement $\delta_{\max}$ for use in the Taylor extrapolation is system dependent. The energy change with atomic displacement depends on the material-specific bond strength, i.e., the force constants of the interatomic bonds. The covalent O–H bonds in water are more rigid than the ionic bonds in LMNTO, so that smaller displacements are needed for water

than for the oxide. As a rule of thumb, we expect the optimal value of $\delta_{\max}$ to be approximately proportional to the smallest interatomic distances (bond lengths) in the reference dataset, which is ~ 0.9 Å for the O–H bond in water and ~ 1.7 Å for O–metal bonds in LMNTO, though the nature of the bonds (covalent, ionic, metallic, or dispersive) will also have an impact.

A related point is the dependence of the energy error on the value of $\delta_{\max}$. As discussed for the oxide system and shown in Fig. 9a, b, the optimal displacement can be a compromise of force and energy accuracy. If the displacement is chosen too large, it is no longer a good approximation that the PES is linear, and the error of the first-order Taylor expansion becomes too large. For each materials system, it is thus necessary to benchmark different displacement values to determine the value that is optimal for both force and energy training.

It is important to note that the Taylor expansion extrapolation approach does not have to be limited to first order. If higher derivatives of the potential energy are available, e.g., if the Hessian matrix has been calculated for the reference structures, then higher-order Taylor expansions can be used to improve the accuracy of the energy extrapolation.

The methodology could be further extended and refined for specific applications. Atomic displacement could be limited to atoms with high force components to avoid introducing noise in shallow regions of the PES. Structures from geometry optimizations are generally not useful for our approach as the atomic forces are near zero, and such structures should not be considered for displacement. It could also be useful to allow selection of specific atoms for displacement, so that the force prediction accuracy can be increased for select substructures. Similarly, the maximal displacement $\delta_{\max}$ could be made a species-specific parameter, which would be especially useful in interface systems containing domains with different types of bonding, such as solid-liquid interfaces. These directions will be explored in future work.

In conclusion, we introduced a new computationally efficient method for the training of accurate artificial neural network (ANN) potentials on interatomic force information and established its effectiveness for different classes of materials. The methodology is based on a Taylor extrapolation of the total energy using the atomic forces from the reference calculations without the need for additional electronic structure calculations. Training occurs on approximate energies, and the computationally demanding evaluation of the second derivatives of the ANN function is not required. Translating the force information to approximate energies makes it possible to bypass the quadratic scaling with the number of atoms that conventional force-training methods exhibit, so that the Taylor-expansion approach can be used for reference datasets containing complex atomic structures. We showed that approximate force training can improve the force prediction error by around 50% of direct force training, and is computationally efficient even for systems for which the latter is challenging or infeasible. We demonstrated for three example systems, a cluster of six water molecules, liquid water, and a complex metal oxide, that the force-training approach (i) allows to substantially reduce the size of the reference datasets for ANN potential construction, by nearly 75% in the case of the water cluster dataset; (ii) increases the transferability of the ANN potential by improving the energy prediction accuracy for unseen structures; and (iii) generally improves the force prediction accuracy, leading to improved stability of MD simulations. The alternative force training approach simplifies the construction of general ANN potentials for the prediction of accurate energies and interatomic forces and is in principle applicable to any type of material.

METHODS

Electronic structure calculations

Water clusters. The reference dataset was generated by performing ab initio molecular dynamics (MD) simulations with DL-POLY⁵⁴ via Chemshell^{55,56}. The reference dataset was generated in an iterative manner. First three MD simulation runs were performed, where the semiempirical GFN-xTB^{48,57} was used as ab initio method. All of these simulations were run for 30 ps with a time step of 0.5 fs. The initial velocities were chosen randomly according to a Maxwell–Boltzmann distribution. The temperature simulated was 300 K for the first two simulations, the third MD simulation run was performed at 800 K. The temperature was in all cases controlled by a Nosé–Hoover thermostat^{58,59}. During the simulations, a harmonic restraint was applied on all atoms to keep the cluster of water molecules confined to a sphere with a radius of about 5 Å. For the restraint a harmonic force with force constant 190.5 eV was applied to any atom that has a distance greater than 0.0005 Å to the 1st atom of the structure (central atom). In order to get a realistic approximation of the energy for the structures obtained from the MD simulation runs, the energies and forces for all structures were recalculated using the BLYP-D3 functional^{60–65} with the basis set def2-TZVP. For the single point energy calculations Turbomole⁶⁶ was used via Chemshell^{55,56}. Using this reference dataset, two neural networks were trained to obtain a first approximation of the ANN potential. On each of the obtained ANN potentials a MD simulation at 300 K was performed for 75 ps. The other parameters for the MD simulation were chosen as discussed before. From these MD simulations on the ANN potentials 4420 additional reference structures were obtained. As for the AIMD reference structures the energies and forces were recalculated on BLYP-D3/def2-TZVP level of theory.

Periodic AIMD simulations. The periodic AIMD simulations were carried out with the Vienna Ab initio simulation package (VASP)^{67,68} and projector-augmented wave (PAW) pseudopotentials⁶⁹. For the bulk water system the revised Perdew–Burke–Ernzerhof density functional⁷⁰ with the Grimme D3 van-der-Waals correction⁶⁴ (revPBE+D3) was used that has previously been shown to be reliable for water²⁶. The AIMD simulations of the Li–Mo–Ni–Ti–O system employed the strongly constrained and appropriately normed (SCAN) semilocal density functional⁷¹.

For both periodic systems, the plane-wave cutoff was 400 eV, and Γ -point only k-point meshes were employed. A time step of 1 fs was used for the integration of the equation of motion, and a Nosé–Hoover thermostat^{58,59} was used to maintain the temperature at 400 K. The bulk water dataset was compiled by collecting every 100th frame from the first 70 ps of the AIMD trajectory, yielding a reference dataset of 700 structures.

Reference data and ANN potential training

The Taylor-expansion approach for force training was implemented in the atomic energy network (aenet) package⁴⁰, which was used for the construction and application of the reported ann potentials. In the present work, the limited memory bfgs algorithm^{72,73} was used for the ann weight optimization (training). The arithr–urban–ceder chebyshev descriptor for local atomic environments²⁰ was employed if not otherwise noted. Further details for the different materials systems follow.

Lennard–Jones dimer. For the dimer dataset, a Chebyshev descriptor with radial expansion order 10 was employed. No angular expansion was used since the dimer does not have any bond angles. The ANN was comprised of two hidden layers with each five nodes and hyperbolic tangent activation, i.e., the ANN architecture was 11–5–5–1.

Water clusters. For the water clusters symmetry functions by Behler and Parrinello¹⁴² were used as descriptor. The parameters for the symmetry functions for water were taken from the publication of T. Morawietz et al.²⁶. For the training of the ANN the reference set was divided randomly into a training and a test set. Thereby 90% of the structures were used as the training set and the remaining 10% of the structures were used as the test set which was used to measure the quality of the predictions obtained from the ANN potential for structures that have not been used in the fit of the potential. The ANN architecture $N_{\text{symm}}-10-10-1$ was used, where N_{symm} is given by the descriptor²⁶ with $N_{\text{symm}} = 27$ for hydrogen and $N_{\text{symm}} = 30$ for oxygen. The hyperbolic tangent was used as activation function.

Each ANN was trained for 5000 epochs, then the weights and biases for the epoch that lead to the smallest test set error during training were used.

Water bulk. For the bulk water ANN potential, a Chebyshev descriptor with a radial expansion order of 18 and an angular expansion order of 4 was used. The interaction cutoffs were 6.5 Å and 3.0 Å for the radial and angular expansion, respectively, and the ANN architecture was 48–10–10–1, i.e., two hidden layers with each ten nodes were used. This corresponds to a total of 611 weight parameters.

Li–Ni–Ti–Mo–O system. A descriptor with an interaction range of 6.5 Å and an expansion order of 20 for the radial distribution function and a range of 3.5 Å and an expansion order of 2 for the angular interactions was used. For the 700-structure dataset, the best balance of model complexity and accuracy was obtained for an ANN potential with two hidden layers and each 10 nodes (48–10–10–1).

ANN potential MD simulations

All ANN potential MD simulations were carried out using the Tinker software⁷⁴ and ANN potentials via an interface with the *ænet* package and used the Verlet algorithm⁷⁵ for the integration of the equation of motion.

The ANN potentials for the water bulk MD simulations were trained on a dataset of ~10,000 frames taken from an AIMD trajectory at 400 K and were based on a 48–20–20–1 ANN architecture. The radial distribution functions (RDF) shown in Fig. 8 were evaluated for a simulation cell with 128 water molecules using MD simulations in the canonical (NVT) ensemble by averaging over a total of 10 ps after a 2 ps equilibration period. A time step of 0.25 fs was used. A Bussi–Parrinello thermostat⁷⁶ was employed for the NVT sampling, and the target temperature was 300 K.

For the MD simulations in Fig. 10, a larger Li–Ni–Ti–Mo–O reference dataset with ~4000 structure was used for training. For this dataset, we employed a 48–20–20–1 ANN architecture. A time-step of 1 fs was used.

DATA AVAILABILITY

The reference datasets of the water cluster, bulk water, and transition-metal oxide systems can be obtained from the Materials Cloud repository (<https://doi.org/10.24435/materialscloud:2020.0037/v1>). These datasets contain atomic structures and interatomic forces in the XCrySDen⁷⁷ structure format (XSF), and total energies are included as additional meta information.

CODE AVAILABILITY

This work made use of the free and open-source atomic energy network (*ænet*) package. The *ænet* source code can be obtained either from the *ænet* website (<http://ann.atomistic.net>) or from [GitHub](#). The Taylor-expansion approach will be made available in the next release of the *ænet* package.

Received: 11 November 2019; Accepted: 10 April 2020;

Published online: 13 May 2020

REFERENCES

- Behler, J. & Parrinello, M. Generalized neural-network representation of high-dimensional potential-energy surfaces. *Phys. Rev. Lett.* **98**, 146401 (2007).
- Bartók, A. P., Payne, M. C., Kondor, R. & Csányi, G. Gaussian approximation potentials: the accuracy of quantum mechanics, without the electrons. *Phys. Rev. Lett.* **104**, 136403 (2010).
- Thompson, A., Swiler, L., Trott, C., Foiles, S. & Tucker, G. Snap: automated generation of quantum-accurate interatomic potentials. *J. Comput. Phys.* **285**, 316–330 (2014).
- Bartók, A. P. & Csányi, G. Gaussian approximation potentials: a brief tutorial introduction. *Int. J. Quantum Chem.* **115**, 1051–1057 (2015).
- Behler, J. First principles neural network potentials for reactive simulations of large molecular and condensed systems. *Angew. Chem. Int. Ed.* **56**, 12828–12840 (2017).
- Zong, H., Pilania, G., Ding, X., Ackland, G. J. & Lookman, T. Developing an interatomic potential for martensitic phase transformations in zirconium by machine learning. *npj Comput. Mater.* **4**, 48 (2018).
- Himanan, L., Geurts, A., Foster, A. S. & Rinke, P. Data-driven materials science: status, challenges, and perspectives. *Adv. Sci.* **6**, 1900808 (2019).
- Nyshadham, C. et al. Machine-learned multi-system surrogate models for materials prediction. *npj Comput. Mater.* **5**, 51 (2019).
- Kostiuchenko, T., Körmann, F., Neugebauer, J. & Shapeev, A. Impact of lattice relaxations on phase transitions in a high-entropy alloy studied by machine-learning potentials. *npj Comput. Mater.* **5**, 55 (2019).
- Deng, Z., Chen, C., Li, X.-G. & Ong, S. P. An electrostatic spectral neighbor analysis potential for lithium nitride. *npj Comput. Mater.* **5**, 75 (2019).
- Schmidt, J., Marques, M. R. G., Botti, S. & Marques, M. A. L. Recent advances and applications of machine learning in solid-state materials science. *npj Comput. Mater.* **5**, 83 (2019).
- Hohenberg, P. & Kohn, W. Inhomogeneous electron gas. *Phys. Rev.* **136**, B864–B871 (1964).
- Kohn, W. & Sham, L. J. Self-consistent equations including exchange and correlation effects. *Phys. Rev.* **140**, A1133–A1138 (1965).
- Burke, K. Perspective on density functional theory. *J. Chem. Phys.* **136**, 150901 (2012).
- Jones, R. O. Density functional theory: its origins, rise to prominence, and future. *Rev. Mod. Phys.* **87**, 897–923 (2015).
- Cooper, A. M. & Kästner, J. Low-temperature kinetic isotope effects in CH₃OH + H₂ shed light on the deuteration of methanol in space. *J. Phys. Chem. A* **123**, 9061–9068 (2019).
- Smith, J. S. et al. Approaching coupled cluster accuracy with a general-purpose neural network potential through transfer learning. *Nat. Commun.* **10**, 2903 (2019).
- Schran, C., Behler, J. & Marx, D. Automated fitting of neural network potentials at coupled cluster accuracy: Protonated water clusters as testing ground. *J. Chem. Theory Comput.* **16**, 88–99 (2019).
- Artrith, N., Morawietz, T. & Behler, J. High-dimensional neural-network potentials for multicomponent systems: applications to zinc oxide. *Phys. Rev. B* **83**, 153101 (2011).
- Artrith, N., Urban, A. & Ceder, G. Efficient and accurate machine-learning interpolation of atomic energies in compositions with many species. *Phys. Rev. B* **96**, 014112 (2017).
- Artrith, N. & Behler, J. High-dimensional neural network potentials for metal surfaces: a prototype study for copper. *Phys. Rev. B* **85**, 045439 (2012).
- Sun, G. & Sautet, P. Metastable structures in cluster catalysis from first-principles: structural ensemble in reaction conditions and metastability triggered reactivity. *J. Am. Chem. Soc.* **140**, 2812–2820 (2018).
- Sosso, G. C., Miceli, G., Caravati, S., Behler, J. & Bernasconi, M. Neural network interatomic potential for the phase change material GeTe. *Phys. Rev. B* **85**, 174103 (2012).
- Artrith, N. & Kolpak, A. M. Understanding the composition and activity of electrocatalytic nanoalloys in aqueous solvents: A combination of DFT and accurate neural network potentials. *Nano Lett.* **14**, 2670–2676 (2014).
- Elias, J. S. et al. Elucidating the nature of the active phase in copper/ceria catalysts for CO oxidation. *ACS Catalysis* **6**, 1675–1679 (2016).
- Morawietz, T., Singraber, A., Dellago, C. & Behler, J. How van der Waals interactions determine the unique properties of water. *Proc. Natl Acad. Sci. USA* **113**, 8368–8373 (2016).
- Smith, J. S., Isayev, O. & Roitberg, A. E. ANI-1: an extensible neural network potential with DFT accuracy at force field computational cost. *Chem. Sci.* **8**, 3192–3203 (2017).
- Cooper, A. M., Hallmen, P. P. & Kästner, J. Potential energy surface interpolation with neural networks for instanton rate calculations. *J. Chem. Phys.* **148**, 094106 (2018).
- Morawietz, T. et al. Hiding in the crowd: spectral signatures of overcoordinated hydrogen-bond environments. *J. Phys. Chem. Lett.* **10**, 6067–6073 (2019).
- Li, W., Ando, Y., Minamitani, E. & Watanabe, S. Study of Li atom diffusion in amorphous Li₃PO₄ with neural network potential. *J. Chem. Phys.* **147**, 214106 (2017).
- Artrith, N., Urban, A. & Ceder, G. Constructing first-principles phase diagrams of amorphous Li_xSi using machine-learning-assisted sampling with an evolutionary algorithm. *J. Chem. Phys.* **148**, 241711 (2018).
- Lacivita, V., Artrith, N. & Ceder, G. The structural and compositional factors that control the Li-ion conductivity in LiPON electrolytes. *Chem. Mater.* **30**, 7077–7090 (2018).
- Artrith, N., Urban, A., Wang, Y. & Ceder, G. Atomic-scale factors that control the rate capability of nanostructured amorphous Si for high-energy-density batteries. *arXiv* <https://arxiv.org/abs/1901.09272> (2019).
- Artrith, N., Hiller, B. & Behler, J. Neural network potentials for metals and oxides—first applications to copper clusters at zinc oxide. *Phys. Status Solidi B* **250**, 1191–1203 (2013).
- Natarajan, S. K. & Behler, J. Neural network molecular dynamics simulations of solid–liquid interfaces: water at low-index copper surfaces. *Phys. Chem. Chem. Phys.* **18**, 28704–28725 (2016).

36. Quaranta, V. et al. Maximally resolved anharmonic OH vibrational spectrum of the water/ZnO(10-10) interface from a high-dimensional neural network potential. *J. Chem. Phys.* **148**, 241720 (2018).
37. Artrith, N. Machine learning for the modeling of interfaces in energy storage and conversion materials. *J. Phys. Energy* **1**, 032002 (2019).
38. Eckhoff, M. & Behler, J. From molecular fragments to the bulk: development of a neural network potential for MOF-5. *J. Chem. Theory Comput.* **15**, 3793–3809 (2019).
39. Singraber, A., Morawietz, T., Behler, J. & Dellago, C. Parallel multistream training of high-dimensional neural network potentials. *J. Chem. Theory Comput.* **15**, 3075–3092 (2019).
40. Artrith, N. & Urban, A. An implementation of artificial neural-network potentials for atomistic materials simulations: Performance for TiO₂. *Comput. Mater. Sci.* **114**, 135–150 (2016).
41. Daw, M. S. & Baskes, M. I. Embedded-atom method: derivation and application to impurities, surfaces, and other defects in metals. *Phys. Rev. B* **29**, 6443–6453 (1984).
42. Behler, J. Atom-centered symmetry functions for constructing high-dimensional neural network potentials. *J. Chem. Phys.* **134**, 074106 (2011).
43. Huang, Y., Kang, J., Goddard, W. A. & Wang, L.-W. Density functional theory based neural network force fields from energy decompositions. *Phys. Rev. B* **99**, 064103 (2019).
44. LeCun, Y., Bottou, L., Orr, G.B. & Müller, K.-R. In *Neural Networks: Tricks of the Trade*, vol. 1524 of *Lecture Notes in Computer Science* (eds. Orr, G. B. & Müller, K.-R.) 9–50 (Springer Berlin Heidelberg, 1998).
45. Lorenz, S. *Reactions on Surfaces with Neural Networks*. Phd thesis, Technischen Universität Berlin, Fakultät II – Mathematik und Naturwissenschaften. <http://opus4.kobv.de/opus4-tuberlin/frontdoor/index/index/docId/202> (2001).
46. Vlack, L., Sun, W. & Kent, P. R. C. Combining configurational energies and forces for molecular force field optimization. *J. Chem. Phys.* **147**, 161713 (2017).
47. Lennard-Jones, J. E. The electronic structure of some diatomic molecules. *Trans. Faraday Soc.* **25**, 668–686 (1929).
48. Grimme, S., Bannwarth, C. & Shushkov, P. A robust and accurate tight-binding quantum chemical method for structures, vibrational frequencies, and non-covalent interactions of large molecular systems parametrized for all spd-block elements ($z = 1-86$). *J. Chem. Theory Comput.* **13**, 1989–2009 (2017).
49. Marsalek, O. & Markland, T. E. Quantum dynamics and spectroscopy of ab initio liquid water: the interplay of nuclear and electronic quantum effects. *J. Phys. Chem. Lett.* **8**, 1545–1551 (2017).
50. Morawietz, T. et al. The interplay of structure and dynamics in the raman spectrum of liquid water over the full frequency and temperature range. *J. Phys. Chem. Lett.* **9**, 851–857 (2018).
51. Lee, J. et al. A new class of high capacity cation-disordered oxides for rechargeable lithium batteries: Li-Ni-Ti-Mo oxides. *Energy Environ. Sci.* **8**, 3255–3265 (2015).
52. Togo, A. & Tanaka, I. First principles phonon calculations in materials science. *Scr. Mater.* **108**, 1–5 (2015).
53. Urban, A., Abdellahi, A., Dacek, S., Artrith, N. & Ceder, G. Electronic-structure origin of cation disorder in transition-metal oxides. *Phys. Rev. Lett.* **119**, 176402 (2017).
54. Todorov, I. T., Smith, W., Trachenko, K. & Dove, M. T. DL-POLY_3: new dimensions in molecular dynamics simulations via massive parallelism. *J. Mater. Chem.* **16**, 1911–1918 (2006).
55. Sherwood, P. et al. QUASI: a general purpose implementation of the QM/MM approach and its application to problems in catalysis. *J. Mol. Struct. THEOCHEM* **632**, 1 (2003).
56. Metz, S., Kästner, J., Sokol, A. A., Keal, T. W. & Sherwood, P. Chemshell—a modular software package for QM/MM simulations. *WIREs Comput. Mol. Sci.* **4**, 101 (2014).
57. Bannwarth, C., Ehlert, S. & Grimme, S. GFN2-xTB—an accurate and broadly parametrized self-consistent tight-binding quantum chemical method with multipole electrostatics and density-dependent dispersion contributions. *J. Chem. Theory Comput.* **15**, 1652–1671 (2019).
58. Nosé, S. A unified formulation of the constant temperature molecular dynamics methods. *J. Chem. Phys.* **81**, 511–519 (1984).
59. Hoover, W. G. Canonical dynamics: equilibrium phase-space distributions. *Phys. Rev. A* **31**, 1695–1697 (1985).
60. Becke, A. D. Density-functional exchange-energy approximation with correct asymptotic behavior. *Phys. Rev. A* **38**, 3098–3100 (1988).
61. Lee, C., Yang, W. & Parr, R. G. Development of the Colle-Salvetti correlation-energy formula into a functional of the electron density. *Phys. Rev. B* **37**, 785–789 (1988).
62. Johnson, B. G., Gill, P. M. W. & Pople, J. A. The performance of a family of density functional methods. *J. Chem. Phys.* **98**, 5612–5626 (1993).
63. Russo, T. V., Martin, R. L. & Hay, P. J. Density functional calculations on first-row transition metals. *J. Chem. Phys.* **101**, 7729–7737 (1994).
64. Grimme, S., Antony, J., Ehrlich, S. & Krieg, H. A consistent and accurate ab initio parametrization of density functional dispersion correction (DFT-D) for the 94 elements H-Pu. *J. Chem. Phys.* **132**, 154104 (2010).
65. Smith, D. G. A., Burns, L. A., Patkowski, K. & Sherrill, C. D. Revised damping parameters for the D3 dispersion correction to density functional theory. *J. Phys. Chem. Lett.* **7**, 2197–2203 (2016).
66. Furche, F. et al. Turbomole. *WIREs Comput. Mol. Sci.* **4**, 91–100 (2014).
67. Kresse, G. & Furthmüller, J. Efficient iterative schemes for ab initio total-energy calculations using a plane-wave basis set. *Phys. Rev. B* **54**, 11169–11186 (1996).
68. Kresse, G. & Furthmüller, J. Efficiency of ab-initio total energy calculations for metals and semiconductors using a plane-wave basis set. *Comput. Mater. Sci.* **6**, 15–50 (1996).
69. Blöchl, P. E. Projector augmented-wave method. *Phys. Rev. B* **50**, 17953–17979 (1994).
70. Zhang, Y. & Yang, W. Comment on *Generalized Gradient Approximation Made Simple*. *Phys. Rev. Lett.* **80**, 890–890 (1998).
71. Sun, J., Ruzsinszky, A. & Perdew, J. Strongly constrained and appropriately normed semilocal density functional. *Phys. Rev. Lett.* **115**, 036402 (2015).
72. Byrd, R., Lu, P., Nocedal, J. & Zhu, C. A limited memory algorithm for bound constrained optimization. *SIAM J. Sci. Comput.* **16**, 1190–1208 (1995).
73. Zhu, C., Byrd, R. H., Lu, P. & Nocedal, J. Algorithm 778: L-BFGS-B: fortran sub-routines for large-scale bound-constrained optimization. *ACM T. Math. Softw.* **23**, 550–560 (1997).
74. Ponder, J. W. & Richards, F. M. An efficient newton-like method for molecular mechanics energy minimization of large molecules. *J. Comput. Chem.* **8**, 1016–1024 (1987).
75. Verlet, L. Computer “experiments” on classical fluids. i. thermodynamical properties of lennard-jones molecules. *Phys. Rev.* **159**, 98–103 (1967).
76. Bussi, G., Donadio, D. & Parrinello, M. Canonical sampling through velocity rescaling. *J. Chem. Phys.* **126**, 014101 (2007).
77. Kokalj, A. XCrySDen—a new program for displaying crystalline structures and electron densities. *J. Mol. Graph. Model.* **17**, 176–179 (1999).

ACKNOWLEDGEMENTS

N.A. and A.U. acknowledge financial support by the U.S. Department of Energy (DOE) Office of Energy Efficiency and Renewable Energy, Vehicle Technologies Office, Contract No. DE-SC0012704. A.C. and J.K. acknowledge financial support by the European Union's Horizon 2020 research and innovation programme (Grant Agreement No. 646717, TUNNELCHEM) and the German Research Foundation (DFG) through the Cluster of Excellence in Simulation Technology (No. EXC 310/2) at the University of Stuttgart. Computational resources for the water cluster DFT calculations were provided by the state of Baden-Württemberg through bwHPC and the German Research Foundation (DFG) through grant no INST 40/467-1 FUGG. Development of *ænet* used the Extreme Science and Engineering Discovery Environment (XSEDE), which is supported by National Science Foundation grant number ACI-1053575. The LMNTO DFT calculations used resources of the Center for Functional Nanomaterials, which is a U.S. DOE Office of Science Facility, at Brookhaven National Laboratory under Contract No. DE-SC0012704. For the simulations of bulk water, we acknowledge computing resources from Columbia University's Shared Research Computing Facility project, which is supported by NIH Research Facility Improvement Grant 1G20RR030893-01, and associated funds from the New York State Empire State Development, Division of Science Technology and Innovation (NYSTAR) Contract C090171, both awarded April 15, 2010. N.A. and A.U. thank Dr. Penghao Xiao for valuable discussions. We thank Dr. Tobias Morawietz for testing our implementation and for insightful discussions.

AUTHOR CONTRIBUTIONS

N.A. designed and planned the study. N.A. and A.U. supervised all aspects of the project. A.C., A.U., and N.A. implemented the methodology into *ænet*, and performed extensive tests. A.C. and N.A. performed DFT calculations and constructed the ANN potentials. A.C., A.U., and N.A. analyzed the results. The manuscript was written by A. C., A.U., and N.A. with the help of J.K. All authors contributed to discussions and commented on the manuscript.

COMPETING INTERESTS

The authors declare no competing interests.

ADDITIONAL INFORMATION

Supplementary information is available for this paper at <https://doi.org/10.1038/s41524-020-0323-8>.

Correspondence and requests for materials should be addressed to N.A.

Reprints and permission information is available at <http://www.nature.com/reprints>

Publisher's note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons license, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons license and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this license, visit <http://creativecommons.org/licenses/by/4.0/>.

© The Author(s) 2020