

ARTICLE OPEN



A systematic approach to generating accurate neural network potentials: the case of carbon

Yusuf Shaidu^{1,2}, Emine Küçükbenli^{1,3}✉, Ruggero Lot¹, Franco Pellegrini⁴, Efthimios Kaxiras^{3,5} and Stefano de Gironcoli^{1,6}✉

Availability of affordable and widely applicable interatomic potentials is the key needed to unlock the riches of modern materials modeling. Artificial neural network-based approaches for generating potentials are promising; however, neural network training requires large amounts of data, sampled adequately from an often unknown potential energy surface. Here we propose a self-consistent approach that is based on crystal structure prediction formalism and is guided by unsupervised data analysis, to construct an accurate, inexpensive, and transferable artificial neural network potential. Using this approach, we construct an interatomic potential for carbon and demonstrate its ability to reproduce first principles results on elastic and vibrational properties for diamond, graphite, and graphene, as well as energy ordering and structural properties of a wide range of crystalline and amorphous phases.

npj Computational Materials (2021)7:52; <https://doi.org/10.1038/s41524-021-00508-6>

INTRODUCTION

The state-of-the-art theoretical framework for computing material properties of crystals at their ground state is density functional theory (DFT)^{1,2}. DFT allows to describe the total energy as a functional of electron density, $E[\rho]$, for a given atomic configuration $\{\mathbf{R}\}$, by taking advantage of the conjugate relationship between the electrostatic potential of the nuclei $V(\{\mathbf{R}\})$, and the ground-state electron density ρ . By solving the expensive quantum mechanical equations that result from this definition for electrons, DFT outlines a path to determine the total energy, the forces on each atom, the stress due to crystal structure, and several other ground-state properties of materials. Yet the cost of solving the quantum mechanical equations, as well as having to work with the extensive electronic wavefunctions and density, hinders the application of this method to systems beyond a few thousands of atoms.

A way to reduce the computational cost lies in the realization that the same conjugate relationship between ρ and V guarantees that a functional exists, which maps the electrostatic potential of the nuclei to the total energy, hence it is possible to describe ground-state properties as a functional of the positions of atoms in the structure, without having to work explicitly with the electron density. Yet, the exact form of such a functional is unknown. One approach to approximate this unknown functional is using artificial neural networks (ANNs). ANNs and in general machine learning techniques have been shown to yield reasonably accurate functional approximations for a wide range of applications, and have already been adopted with success to some material science problems^{3–15}.

ANNs can be seen as an attractive alternative to the classical approach for constructing interatomic interaction models (also known as force fields (FFs)) where physical intuition is used to fix the form of the approximate functional for $E[V(\{\mathbf{R}\})]$. While physically meaningful forms can describe the interatomic interaction in a compact way, with only few parameters to be fitted, the rigidity of the functional form reduces the predictive power of this method in exploratory studies. In particular, for highly

polymorphic materials such as carbon, where several different bonding types and structures exist, the lack of transferability of a model from one structure to another results in many different interaction models, each with a limited applicability. For example, among the several empirical FFs for carbon, the non-reactive, short range, bond-order-based Tersoff¹⁶ model can describe dense sp^3 carbon structures while a highly parametric reactive force field (ReaxFF)¹⁷ that explicitly includes long-range van der Waals (vdW) interactions and Coulomb energy through charge equilibration scheme¹⁸ is needed for structures with sp^2 hybridization. Furthermore, even though these empirical FFs give a qualitative understanding of materials properties, they are quantitatively inaccurate when compared to both ab initio methods and experiments^{19–22}.

Interatomic interaction models based on ANNs do not have a fixed functional form beyond the network architecture, and their parameters are fitted to vast amounts of ab initio quantum mechanical data in the hope of assimilating the physics of the system into the parametrization. Hence the transferability restraint of classical FFs, that is due to their rigid form, is traded for a transferability challenge in the case of neural networks due to the (lack of) variety and completeness in the training set. To address this challenge of generating truly transferable ANN interatomic interaction models, training data must be obtained from an efficient and thorough sampling of the potential energy landscape. Such sampling of the very rugged and high dimensional landscape with ab initio electronic structure tools is a formidable challenge.

In this work, we integrate evolutionary algorithm (EA) with molecular dynamics (MD) and clustering techniques in a self-consistent manner to sample the potential energy landscape and obtain data with high variability. The workflow we introduce extends the training data iteratively, similar to other active learning approaches that previously appeared in literature^{19,23–26}. Unlike these methods that aim at constructing an optimal dataset for a specified part of the potential energy landscape, our workflow targets an unbiased training dataset, which is necessary for increased transferability expected of a general purpose

¹Scuola Internazionale Superiore di Studi Avanzati, Trieste, Italy. ²The Abdus Salam International Centre for Theoretical Physics, Trieste, Italy. ³John A. Paulson School of Engineering and Applied Sciences, Harvard University, Cambridge, MA, USA. ⁴Laboratoire de Physique de l'École Normale Supérieure, Université PSL, CNRS, Sorbonne Université, Université de Paris, Paris, France. ⁵Department of Physics, Harvard University, Cambridge, MA, USA. ⁶CNR-IOM DEMOCRITOS, Istituto Officina dei Materiali, Trieste, Italy. ✉email: kucukben@g.harvard.edu; degironc@sissa.it

potential. Moreover, for reliable materials modeling, it is crucial to have indicators that signal when the limit of transferability is crossed. We address this aspect of ANN models by studying the relationship between data variability and transferability of the trained network via unsupervised data analysis. We demonstrate the performance of the approach highlighted above on the challenging example of crystalline and amorphous carbon structures.

This study is a continuation of similar efforts in the literature: the first ANN interaction model for elemental carbon was developed in 2010 by Khaliullin et al.¹⁹ to study graphite–diamond co-existence. The network was trained on an adaptive training set, where the starting configurations were manually selected from randomly distorted graphite and diamond phases, relaxed under a range of external pressures (from -10 to 200 GPa) at zero temperature. Then, configurations for new training data were obtained using this model in finite temperature MD simulations, which in turn were used to refine the network, until a self-consistency was reached in the prediction error on the new structures. More recently in 2019, a hybrid model, where an ANN potential for the short-range interaction is supplemented with a theoretically motivated analytical term to model long-range dispersion, has been developed in order to address the properties of monolayer and multilayer graphene, with encouraging results²². As we will demonstrate in this work, ANN models such as these, built on data sampled solely from a limited part of the potential energy landscape can, however, be highly non-transferable. This transferability challenge for carbon has been observed with kernel-based machine learning models as well.

In 2017, a kernel-based model, specifically, a Gaussian approximation potential (GAP), was constructed²¹ using data from MD melt-quench trajectories of liquid and amorphous carbon, to study amorphous structures. Motivated from its non-optimal behavior on crystalline phases, authors developed

another GAP model with a specialized training data obtained via MD, for graphene²⁷. It is worthwhile to note that recently, a strategy combining kernel-based model generation with crystal structure prediction was suggested by Bernstein et al.²⁸. Since computational cost for training or evaluation of a kernel-based model grows with the training set, however, this approach is suitable for small scale configuration space sampling. Alternatively, a sparsification approach, such as the one based on clustering recently proposed in ref.²⁹, can be used. In comparison, computational cost of neural networks is independent of the size of the training dataset, a feature that is exploited in the current study for accurate prediction of elastic and vibrational properties. It should be mentioned that regression-based machine-learned potential models other than GAP also exist, e.g., spectral neighbor analysis potential (SNAP)⁸ and moment tensor potential (MTP)³⁰. A recent work comparing them concludes GAP to have the highest accuracy, but also the highest computational cost, increasing with the size of the training dataset³¹. SNAP and MTP use lower cost regression strategies to correlate the local atomic environment with its contribution to the total energy.

In this work we use a systematic approach to construct a highly flexible and transferable neural network potential (NNP) and demonstrate its application to the development of a general NNP for carbon. We compare its performance with respect to other potential models previously optimized for specific phases and discuss the implications of our results for the trade-off between transferability and specialization.

RESULTS

Self-consistent training and validation

The NNP is constructed following the self-consistent approach sketched in Fig. 1. This recursive data-creating and fitting cycle starts with a trial FF, which is used to generate an initial set of

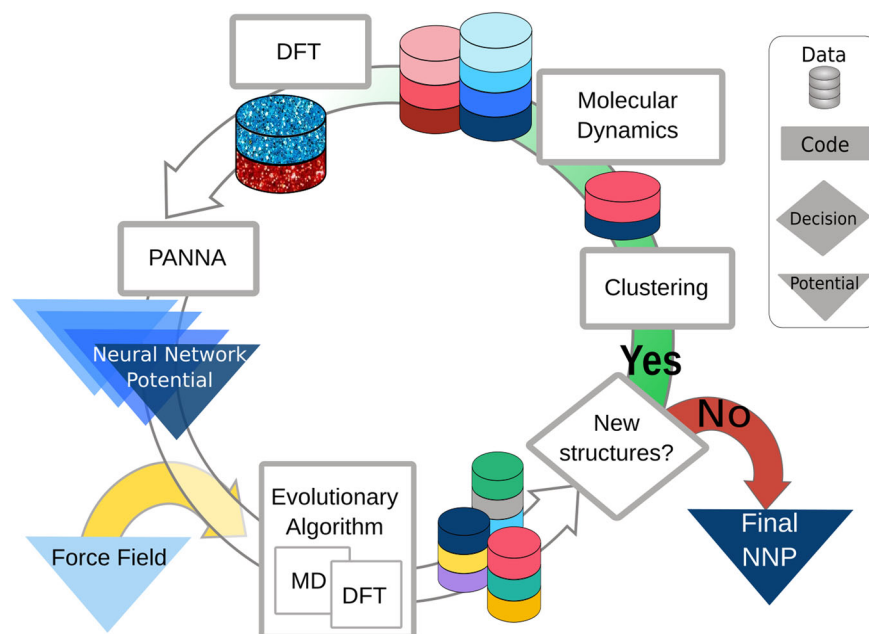


Fig. 1 The self-consistent scheme. The initial step to start the process (yellow arrow) can be performed with a classical force field as shown here, or any comprehensive dataset of structures such as the ones in Aflowlib⁷², Materials Genome Initiative⁷³, or Nomad⁷⁴ repositories can be used to generate the first neural network potential model (blue triangle) to be refined through the self-consistent cycle. Once an initial potential model is chosen, evolutionary algorithm enables a diverse set of structures to be sampled. The following clustering-based pruning of structures further ensures that no single polymorph biases the dataset, i.e., at each step only novel structures (red and blue disks for the particular step highlighted above) are to be considered, further refined, and added to the dataset. The subsequent MD simulations sample the potential energy surface of each polymorph. Finally, DFT calculations performed on a subset of MD-sampled structures are added to the ab initio dataset obtained thus far. The ab initio dataset augmented this way is then used to train the next neural network potential model (a darker blue triangle), starting the next cycle of the self-consistent scheme until no new structures are found by the evolutionary algorithm.

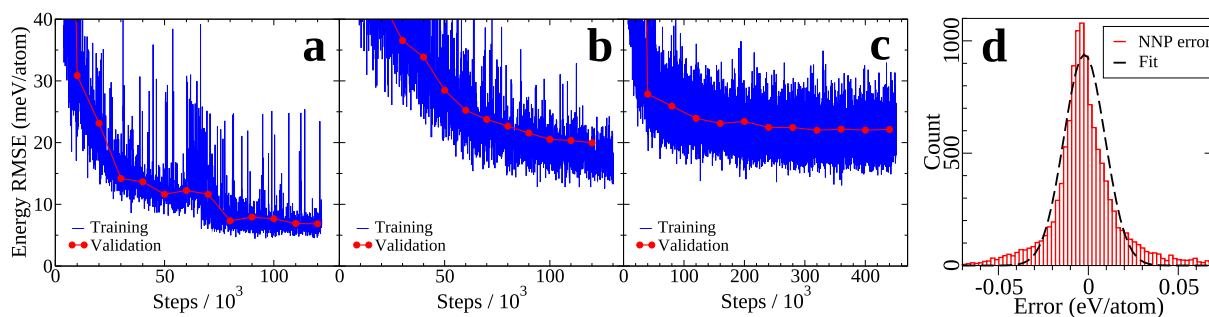


Fig. 2 The evolution of the distribution of error in energy prediction. The RMSE in prediction of per atom energy for potentials trained at first, second, and third iteration of the self-consistent cycle is given in **a–c**, respectively. The blue lines are the RMSE on a given batch of 128 configurations during training. The networks are evaluated during training on all the validation set of sizes ≈ 3000 , ≈ 5200 and, ≈ 12000 configurations for first, second, and third iterations, respectively (red dots with lines as guide to eye). The final training and validation RMSE are reported in Table 1. **d** Error distribution for the validation dataset at third iteration. The black dashed line is a normalized Gaussian fit, resulting in an RMSE of 11.3 meV, clearly failing to fit the fat-tailed distribution. The error distribution of energy and force prediction for all iterations is given in Supplementary Fig. 3.

configurations via EA. In the absence of an established FF model for a new material, rough approximations such as Lennard–Jones or low-cost DFT approximations can be used with small unit cells for the very first iteration. EAs are commonly used in crystal structure prediction studies as they allow efficient sampling of the configuration space. Their success in thorough sampling is demonstrated by their ability to predict new crystal structures before the experimental observation^{32,33}. As the exploration of the configuration space continues, a single-point DFT calculation is performed on each distinct polymorph generated by EA. These structures are then clustered using a distance measure. From each cluster, a representative example is manually selected and a classical MD simulation at a given pressure and temperature range is performed. The additional MD simulation step allows the sampling of the whole neighborhood of the equilibrium configuration for each polymorph, resulting in accurate prediction of structural properties for every polymorph. The dataset obtained this way is used to train a neural network model. The trained NNP is then used for starting a new iteration of the self-consistent cycle. This increases the training set diversity, by preventing the energetically favorable structures that are easily accessed by EA from dominating the whole training set. The iterative procedure highlighted above is repeated until no new structures are found.

While iterative expansion of training set is not a new idea, our implementation pushes its limits in diversity and balance: we use a full EA to sample configurations, without anchoring the search in any known polymorph or rigid transformations between polymorphs as in refs. 25 or 26. This makes our method applicable to materials with unexplored phase space and prevents any bias toward known phases. We then use clustering, which allows to achieve a balanced set despite the tendency of EA to sample stable configurations more often. Finally, starting from a representative configuration for each cluster, we perform MD simulations so that equilibrium properties of every polymorph are well described independent of their stability with respect to the ground state. We refrain from using active learning methods that depend on network agreement (as in ref. 23) as network prediction errors are not guaranteed to be uncorrelated, e.g., two networks may agree on the wrong result, especially if under-parametrized. We also refrain from expanding the training set with structures obtained solely through MD trajectories as in ref. 34, because of the risk of missing significant polymorphs that would only be sampled rarely, and with decreasing frequency, i.e., requiring longer and longer MD runs to run into significant additions to the dataset. Instead, a coherent integration of EA, clustering and MD together yields an unbiased, balanced, and diverse dataset. Further details of the self-consistent training used in this work

Table 1. Training and validation RMSE.

Iteration	Dataset size	$\epsilon_E^{\text{train}}$	ϵ_E^{val}	$\epsilon_F^{\text{train}}$	ϵ_F^{val}
1	15,841	6.8	6.8	0.14	0.14
2	30,815	17.1	20.0	0.19	0.22
3	60,133	22.0	22.1	0.26	0.27

For each iteration we report the RMSE in energy on training ($\epsilon_E^{\text{train}}$) and validation set (ϵ_E^{val}) in eV/atom, and the RMSE in the cartesian components of the force vector for training ($\epsilon_F^{\text{train}}$) and validation set (ϵ_F^{val}) in eV \AA^{-1} . The training RMSE is the average over the batch RMSE of the last 2500 training steps, while the validation RMSE is evaluated over the entire validation set with the NNP obtained at the last training step. The energy (force) mean absolute error on the validation set is 4 (0.09), 12(0.12), and 14(0.16) meV/atom (eV \AA^{-1}) for first, second, and third iteration, respectively.

are given in “Methods” and the expansion of the dataset explored at each step is given in the Supplementary Fig. 1.

The performance of an NNP at each self-consistent loop is evaluated during training via the validation scheme. Figure 2 shows the evolution of NNP energy accuracy on the training and validation set as a function of training steps at each self-consistent iteration (Fig. 2a–c). The training root-mean-square error (RMSE) corresponds to the instantaneous RMSE computed on the elements of the batch considered at that training step while the validation RMSE is computed on all the configurations in the validation set. The RMSE on the validation set agrees with the training RMSE throughout the training, an indication that the model does not overfit to the training dataset. The analysis of the force prediction error at different stages of training gives similar results and can be found in the Supplementary Fig. 2. The increase in energy and force RMSE from iteration 1 to 3 is a result of the increase in the diversity of atomic environments. At each self-consistent iteration, the diversity of the dataset increases as new structures are explored (see Table 1), while the number of parameters of the network, therefore its capacity, is kept fixed. It is worth noting that the prediction error is not distributed according to a Gaussian distribution function but a fatter-tailed one (see Fig. 2d). Therefore, while the RMSE given here is a good measure to compare training and validation error with one another, it overestimates the average NNP prediction error in general.

To demonstrate how the general accuracy of the NNPs is changing with each iteration, we check their performance on a dataset of 197 distinct carbon structures. These structures were obtained by Deringer and co-workers³⁵ via random search of crystal structure of carbon with a GAP developed for liquid and amorphous carbon systems²¹ and are distributed online³⁶. They

represent 197 different crystal configurations of carbon, classified according to the topology of the carbon network. For consistency, their energies are re-calculated with the same DFT parameters as explained in “Methods”. Figure 3 shows the energy ranking as predicted by NNP, GAP, Tersoff, and ReaxFF. It can be seen that the NNP accuracy gets better with each iteration. The third iteration NNP accuracy agrees remarkably well with DFT results and performs better than all the other methods tested. It is noteworthy that the final NNP carries no signature of the ReaxFF used in the initial step to explore the configuration space. Both classical potentials, Tersoff and ReaxFF, perform very poorly compared to machine-learned ones, and the NNP outperforms GAP results published in refs. ^{21,35}, albeit GAP was fitted on ab initio data obtained with local




Fig. 3 Prediction of energy ordering of carbon structures. The energy ordering predicted by NNP is compared to DFT and other models for 197 distinct carbon structures reported in ref. ³⁵. **a** Prediction performance of NNP at different iterations of the self-consistent cycle visible improves. **b** Prediction performance of GAP²¹, reactive force field (ReaxFF)⁷⁵ and Tersoff¹⁶ models is reported alongside the final NNP model (blue line). For comparison, we train a new model with LDA exchange-correlation functional, named as NNP-LDA (red line). The neural network potentials for the two functionals overlap for the majority of the structures, as do the DFT results (see the Supplementary Fig. 4). Further analysis of prediction error in energy (instead of ranking) and the analysis of similarity between this dataset and the one used in NNP generation is reported in Supplementary Fig. 5.

density approximation (LDA) exchange-correlation functional³⁷. For fair comparison, we train a new NNP, using the same training dataset structures obtained via the self-consistent procedure, but using LDA functional. This potential, referred as NNP-LDA, performs similarly to the NNP highlighted in this work, and similarly outperforms all the other potentials. In the rest of the work, the results denoted with NNP refer to the potential that is trained with the rVV10 functional unless otherwise specified.

Structural and elastic properties

In this section, we discuss the performance of the NNP on the structural and elastic properties of select carbon polymorphs, namely, diamond, graphite, and graphene (see Tables 2–4). The equilibrium lattice parameters are obtained by minimizing the total energy until the force components on each atom are lower than $26 \text{ meV } \text{Å}^{-1}$ for both DFT and NNP simulations. We also include results obtained with Tersoff potential, as well as other DFT and machine learning studies in literature.

In the case of diamond, all machine learning methods agree reasonably well with the DFT results they were trained with, both for the equilibrium volume and elastic constants. The largest deviation is seen in C_{12} prediction with GAP with 24% relative error. For all properties tested, the predictions of NNP of the current study is within a relative error of 5% with respect to DFT. It should be noted that the variation between DFT studies employing different exchange-correlation functionals is larger than the difference between machine-learned models and the DFT results they are trained to reproduce. Tersoff potential, although it predicts the equilibrium volume well, fails to predict the C_{44} .

In the more challenging case of graphite, C_{11} and C_{12} relate to the in-plane elastic properties while C_{33} probes the relationship between strain and stress between the planes, which are held together by vdW interactions. C_{13} and C_{44} couple the strong in-plane interaction with the weak out-of-plane ones, namely C_{13} can be seen as a measure of interlayer dilation upon layer compression, and C_{44} as a measure of response to shear deformation. The performance of the NNP on prediction of graphite elastic constants is aligned with this overview: for all potentials reported in Table 3, in-plane lattice parameter and elastic constants are better predicted than the ones that relate to out-of-plane interaction, indicating that more data or better training is needed to describe these more delicate properties. Yet it is encouraging that the general purpose NNP of the current work performs at least as well as other NNPs from literature that were developed with a focus on vdW systems such as graphite and multilayer graphene. In the “Discussion,” we discuss how focusing on particular system could further improve on these predictions.

Table 2. Elastic properties of diamond.

Property	NNP	DFT	GAP ^a	DFT ^a	NNP ^b	DFT ^b	Tersoff ¹⁶	Exp ⁶⁴	Exp ⁶⁵
a (Å)	3.576	3.584	3.539	3.532	3.569	3.570	3.566	–	3.567
B_0 (GPa)	431	425	438	466	434	439	426	442	445
C_{11} (GPa)	1054	1044	1090	1101	1016	1056	1074	1079 (5)	1080
C_{12} (GPa)	119	116	112	148	142	130	102	124 (5)	127
C_{44} (GPa)	542	547	594	592	580	567	641	578 (2)	576

All machine-learned potentials reproduce their reference DFT lattice parameter with <1% relative error. When the comparison extends to elastic constants and bulk modulus, however, only the NNP described in this work shows a consistently close agreement between DFT and the potential model, <3% relative error, the range of variation also observed between two different experiments. Energy as a function of volume is given in the Supplementary Fig. 6.

^aReference²¹.

^bReference¹⁹.

Table 3. Elastic properties of graphite.

Property	NNP	DFT	NNP ^a	DFT ^a	hNN-Gr _x ^b	DFT ^b	Experiment
a (Å)	2.471	2.471	2.467	2.467	2.467	2.466	2.464 ^c , 2.463 ^d
c (Å)	6.732	6.719	6.688	6.815	6.804	6.800	6.712 ^c , 6.712 ^d
B_0 (GPa)	48	40	48	37	–	–	36 (11) ^c
C_{11} (GPa)	1053	1048	1080	1069	978	1080	1060 (20) ^e , 1109 (16) ^c
C_{12} (GPa)	197	182	179	162	177	162	180 (20) ^e , 139 (36) ^c
C_{13} (GPa)	–23	–5	0	–4	–67	–5	15 (5) ^e , 0 (3) ^c
C_{33} (GPa)	57	43	52	40	40	33	37 (10) ^e , 39 (7) ^c
C_{44} (GPa)	–5	4	7	5	1.79	3.36	0.27 ^e , 5 (3) ^c
C_{66} (GPa)	428	433	–	–	–	–	485 (11) ^c

Here we report the lower bound for the bulk modulus using Reuss average, i.e., $1/B_0 \equiv s_{11} + s_{22} + s_{33} + 2(s_{12} + s_{23} + s_{31})$. The robust intraplanar structural features of graphite is captured well by all machine-learned potentials while the weaker interplanar interaction and, in particular, elastic properties that couple the two, are more challenging to capture. This is true even for the hybrid potential hNN-Gr of ref. ²² where the distance dependence of the long-range interaction is manually set to r^{-6} and the potential is tailor-fit to describe multi-graphene systems. Differences between DFT references are of similar magnitude as the differences between NNP and DFT. Energy as a function of volume is given in the Supplementary Fig. 6.

^aReference ¹⁹.

^bReference ²².

^cReference ⁶⁶.

^dReference ⁶⁷.

^eReference ⁶⁸.

Table 4. Elastic properties of graphene.

Property	NNP	DFT	hNN-Gr _x ^a	DFT ^a	Tersoff ¹⁶	Experiment
a (Å)	2.470	2.470	2.467	2.466	2.530	2.46 ^b
ν	0.244	0.173	0.197	0.149	–0.158	–
E (GPa)	967	1015	1021	1060	1216	1015(149) ^c , 2400 (400) ^d
C_{11} (GPa)	1028	1047	1062	1084	1247	–
C_{12} (GPa)	251	181	209	161	–197	–

The graphene 2D elastic constants were computed with the normalized 3D stress as $\sigma_{2D}(\epsilon) = E(\epsilon)/A_0$ at a given strain of ϵ , where $E(\epsilon)$ is the total energy at ϵ and $A_0 = \sqrt{3}a^2/2$ is the area of graphene plane. E is the Young modulus and ν is Poisson's ratio. The elastic constant is converted to bulk properties in GPa by dividing by the interlayer distance $c/2$ of graphite reported in Table 3 or in respective experimental reference. Both machine-learned potentials similarly overestimate the Poisson's ratio and underestimate the Young's modulus with respect to their DFT references. Differences between DFT references are of similar magnitude as the differences between NNP and DFT. Energy as a function of volume is given in the Supplementary Fig. 6.

^aReference ²².

^bReference ⁶⁹.

^cReference ⁷⁰.

^dReference ⁷¹.

Vibrational properties

Phonon dispersion relations give a complete picture of the elastic properties of a material, and reproduction of the dispersion relations obtained via DFT is a tight accuracy criterion on model potentials. Here we examine the performance of NNP through its prediction of phonon dispersion in the case of diamond and graphene, as a function of lattice parameter, up to a 1% deviation from the equilibrium structure. This is a relevant range for thermal

expansion of these materials as, for instance, the change in lattice parameter of diamond at temperatures up to 2000 K is found to be below 1%³⁸. Similarly, thermal expansion increases graphene lattice parameter only within 1% at temperatures up to 2500 K³⁹.

The predictions of NNP for phonon dispersion of diamond and graphene are depicted in Fig. 4. There is an overall good agreement between NNP and DFT in the case of diamond. In the case of graphene, there is a slight disagreement for the transverse optical mode around K point. This is the same trend observed in other machine-learned potentials^{22,27} and likely the result of electronic structural properties associated with this special point coupling with the lattice vibration. For both structures, the predicted phonon frequencies reduce when the crystal expands and increase when it is compressed, as expected. An exception to this is the soft flexural mode of graphene close to Γ point. The instability of graphene upon compression can be seen via small imaginary frequency of this mode (shown as negative). This feature is predicted with DFT and is successfully reproduced with NNP, pointing at the capacity of NNP in predicting important structural stability indicators.

Phonon dispersion of graphite, shown in Fig. 5 displays negative frequencies for low wave vectors close to Γ , along the perpendicular direction to the graphene plane. These phonon modes are particularly soft and are very sensitive to the level of accuracy of the forces predicted by NNP. We verify this hypothesis with an alternative loss function for NNP training, one that minimizes the relative force error rather than the absolute one used so far (see “Methods”). With a loss function that is based on relative error, configurations with small forces impact the NNP parameter minimization more strongly. We retrain the NNP starting from the previously optimized parameters and report graphite phonon dispersion obtained with the retrained NNP in Fig. 5b. It is evident that this approach can improve the NNP prediction for structures with small forces, e.g., close to equilibrium conditions. Phonon dispersions for diamond and graphene obtained with this NNP are given in Supplementary Fig. 7, and demonstrate that the general quality of the NNP is slightly modified and mostly for the high frequency modes. Further tuning of retraining parameters and loss function can be used as a way to achieve higher accuracy in the desired range of energy and force distributions.

An alternative approach that is commonly used in literature for improving NNP prediction is to bias the training set with the configurations for a certain polymorph. To show the effect of this approach, we train the NNP model from scratch this time using a biased dataset with structures from the close neighborhoods of diamond or graphite only. The results reported in Fig. 5c show that this approach indeed allows to reach a better agreement with DFT and there are no imaginary phonon frequencies. However, as it will be further examined later (see Discussion), while this NNP model predicts well properties of configurations around its reference, i.e., diamond or graphite, it is found to be highly non-transferable to other regions of the potential energy surface of carbon.

Amorphous carbon structures

Last, we test the NNP in its ability to construct amorphous carbon structures in a range of densities from 1.5 to 3.5 g cm^{–3} generated via the melt and quench method following the steps highlighted in ref. ²¹. We start from a 216 atoms simple-cubic simulation cell and randomized velocities at 9000 K and perform MD simulation first at 9000 K with Nose–Hoover thermostat⁴⁰ for 4 ps, followed by another at 5000 K for 4 ps, then a fast exponential quench to 300 K at a rate of 10 K fs^{–1} (total duration ~0.5 ps), and finally for 4 ps we let the system evolve with the thermostat fixed at 300 K.

The radial distribution function (RDF) of liquid and amorphous phases are given in Fig. 6a. The liquid is less ordered than the

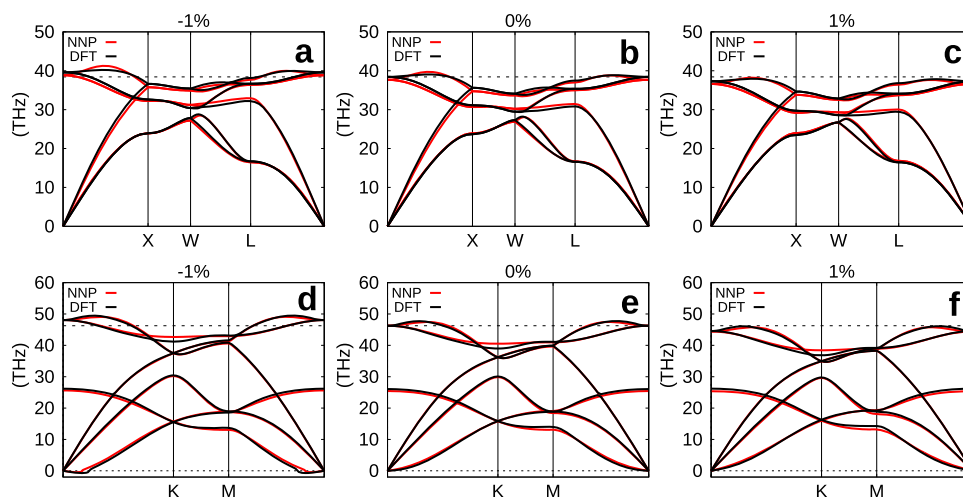


Fig. 4 Phonon dispersion at equilibrium and deformed geometries. The phonon dispersion along the high symmetry lines of diamond and graphene is reported in **a–c** and **d–f**, respectively. The value at the top of each graph represent the percentage of expansion (positive) or compression (negative) of the lattice parameter. The dotted black line is the maximum frequency in THz at the Γ point at equilibrium lattice parameter.

amorphous configurations at all densities, for all potentials considered. In ref.²¹, it was shown that both DFT and GAP have a non-zero first minimum for the liquid phase at about 1.9 Å, which is not properly described by the screened Tersoff potential⁴¹. Similarly, the NNP of this work captures the non-zero first minimum in the liquid phase while the original Tersoff potential does not. In the case of the amorphous phase, historically one of the first validation cases for the Tersoff potential, the agreement is overall better. A more detailed comparison of RDF reported in ref.²¹ and experiments is given in the Supplementary Fig. 8 and shows that NNP can successfully reproduce peak position and width across the densities considered.

In order to quantify the short-range order of amorphous structures, we calculate the sp^3 concentration by computing the fraction of carbon atoms with at least four neighbors within a 1.85 Å radius. In Fig. 6b, we show the behavior of this quantity as a function of density, comparing with the results of ref.²¹ and those obtained with regular and screened Tersoff potentials⁴¹. All methods underestimate the experimental observations yet show a similar general trend with density.

There are quantitative differences among the predictions of theoretical models, in particular, the difference between NNP and GAP predictions are more significant at medium and low densities. This may be attributed to the fact that the DFT dataset used to construct the GAP potential is built with LDA, while in this study the DFT dataset for NNP is built with an accurate exchange-correlation functional that includes vdW interaction from first principles. In the low density region, vdW interactions allow bonding beyond the typical sp^3 bond length, such that low energy configurations can be constructed with less sp^3 and more sp^2 bonds; while at high densities and at shorter length scales, vdW interactions are of lesser significance. This is more evident as we compare the sp^3 count predicted with NNP-LDA as it agrees more closely with the GAP result, revealing the role of the underlying DFT reference in the prediction of the properties of amorphous materials with machine-learned potential models.

The bonding character between atoms strongly affects the elastic properties of materials. Hence, comparing the elastic properties as observed by experiments with those predicted by theory is another way of assessing the theoretical prediction of sp^3 count in amorphous structures. In order to do that, we first find the metastable configurations closest in the phase space to the amorphous structures examined so far, by further quenching the

dynamics from 300 to 0K, and then performing geometry relaxation until the force components on atoms are below 1 mRy bohr⁻¹ at fixed volume. Figure 6c shows the Young's modulus of these metastable amorphous structures as a function of density. The agreement with the experiment is remarkable, hinting that the discrepancy in theoretical and experimental sp^3 count seen in Fig. 6b might stem from an inconsistency in definitions between theory and experiment, i.e., the neighbor count within 1.85 Å used in theory underestimates the experimentally measured value that is obtained via comparison of electron energy-loss spectroscopy peak area to graphitized carbon^{42,43}.

We emphasize that the NNP was not constructed specifically for the description of amorphous C, nor did it include amorphous or melt structures hand-picked to represent these configurations. Despite this, the self-consistent approach yields an NNP, which describes these structures well at all volumes considered, validating successful extrapolation of the potential beyond the training set (see Supplementary Fig. 9 for energy analysis of liquid and amorphous structures compared to the training set).

DISCUSSION

The accuracy of a neural network model is often measured by the distribution of the prediction error on a test dataset, in particular via mean and standard deviation of error. But as is the case with training sets, test sets are also not standardized between studies. Therefore the accuracy of potentials tested on different datasets cannot be compared. Here we study the effect of the training and test sets on the apparent accuracy of networks, and measure the impact of these sets on the transferability of NNPs.

For every configuration in a dataset, we first define its Euclidean distance from a reference atomic environment (e.g., cubic diamond, graphite). The distance between the reference configuration α and a given configuration β is defined as:

$$d_{\alpha\beta} = \frac{1}{2} \left(\frac{1}{N_{\beta}^{\text{at}}} \sum_{i=1}^{N_{\beta}^{\text{at}}} |\mathbf{g}_{\alpha} - \mathbf{g}_{\beta}^i|^2 \right)^{1/2} \quad (1)$$

where $\mathbf{g} = \frac{\mathbf{g}}{|\mathbf{g}|}$ with \mathbf{G} being a fingerprint vector that describes the atomic environment of all atoms in the unit cell for a given configuration, N_{β}^{at} is the number of atoms in configuration β . In this work, for the definition of atomic environment, we use the well-established atom-centered symmetry functions of Behler and

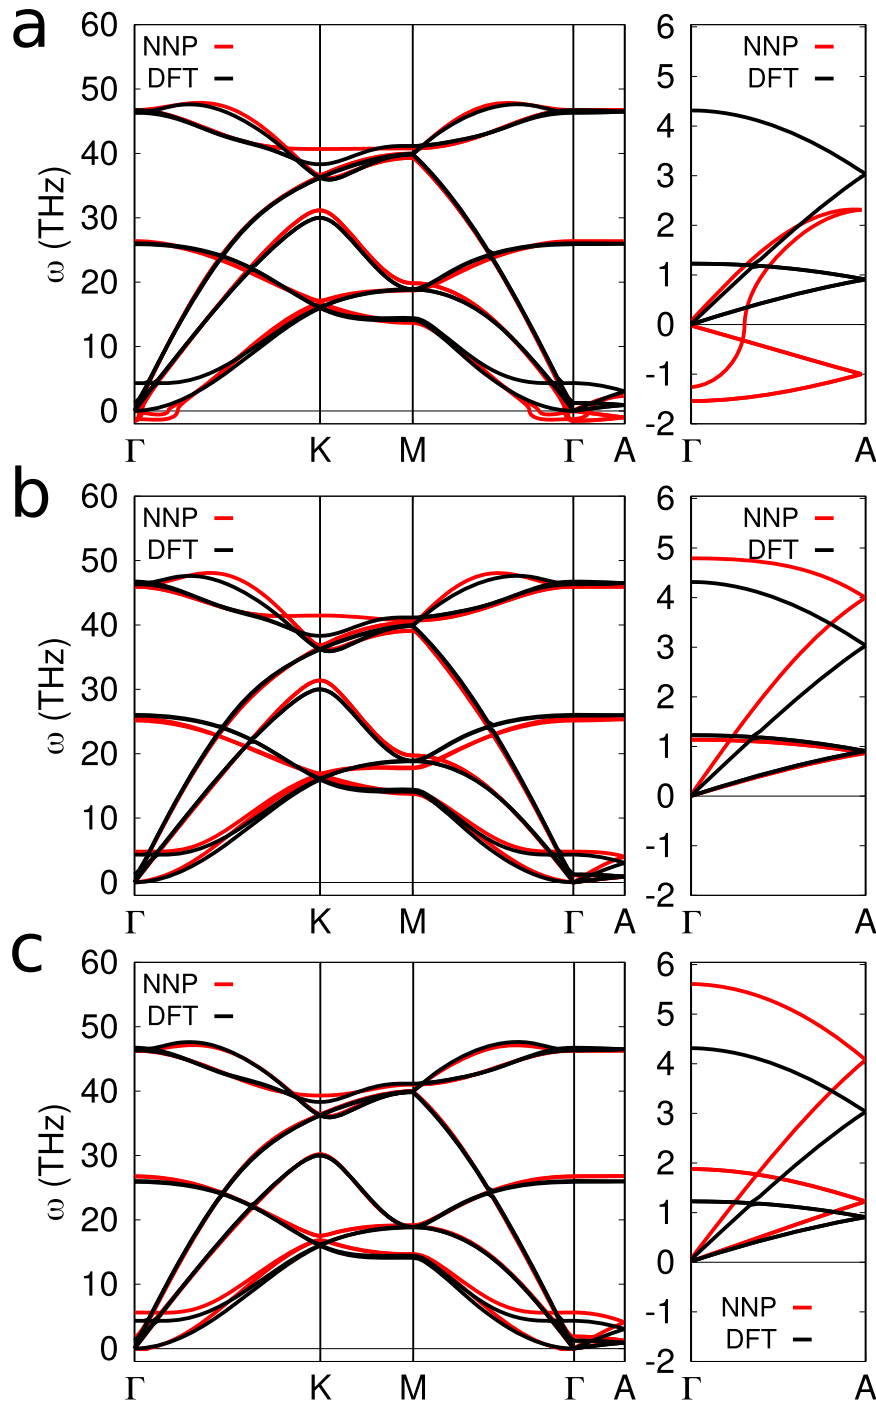


Fig. 5 Phonon dispersion of graphite. The phonon dispersion along the high symmetry lines is reported for **a** an NNP trained with the whole dataset at the last iteration, **b** an NNP retrained with the whole dataset but with the minimization of the relative error on forces, and **c** an NNP trained with all the data within $D = 0.05$ from diamond and graphite (D_{12} , as described in “Discussion”). The small imaginary frequencies are lifted by modifying the NNP training loss function, or by training on data close to graphite in structure.

Parrinello⁴⁴, with modifications by refs. ^{45,46}. This definition is also used to describe the input to the neural network architecture. (see “Methods” for a detail description of the descriptor vectors and their use in neural network training.)

Then, we construct a dataset by considering only configurations within a given cutoff distance D from this reference. Following this strategy we build four datasets, three of which are referenced from cubic diamond with D values of 0.05, 0.10, and 0.15; the fourth one is referenced from either cubic diamond or graphite with $D = 0.05$ (denoted by D_{12}). For each D , 20% of the dataset is

set aside for validation and the remaining 80% is used for training. We train four different NNPs on these four sets from scratch, and test each on the respective validation datasets.

In Fig. 7a, we report the training and validation RMSE in energy prediction as the cutoff distance D from the reference structure increases. We show that an RMSE as low as 2.4 (2.5) meV/atom for training (validation) can be obtained when training and validation configurations are very similar, i.e., within a distance of 0.05 from the diamond reference. However, the prediction error of this NNP dramatically increases as it gets tested on structures farther in the

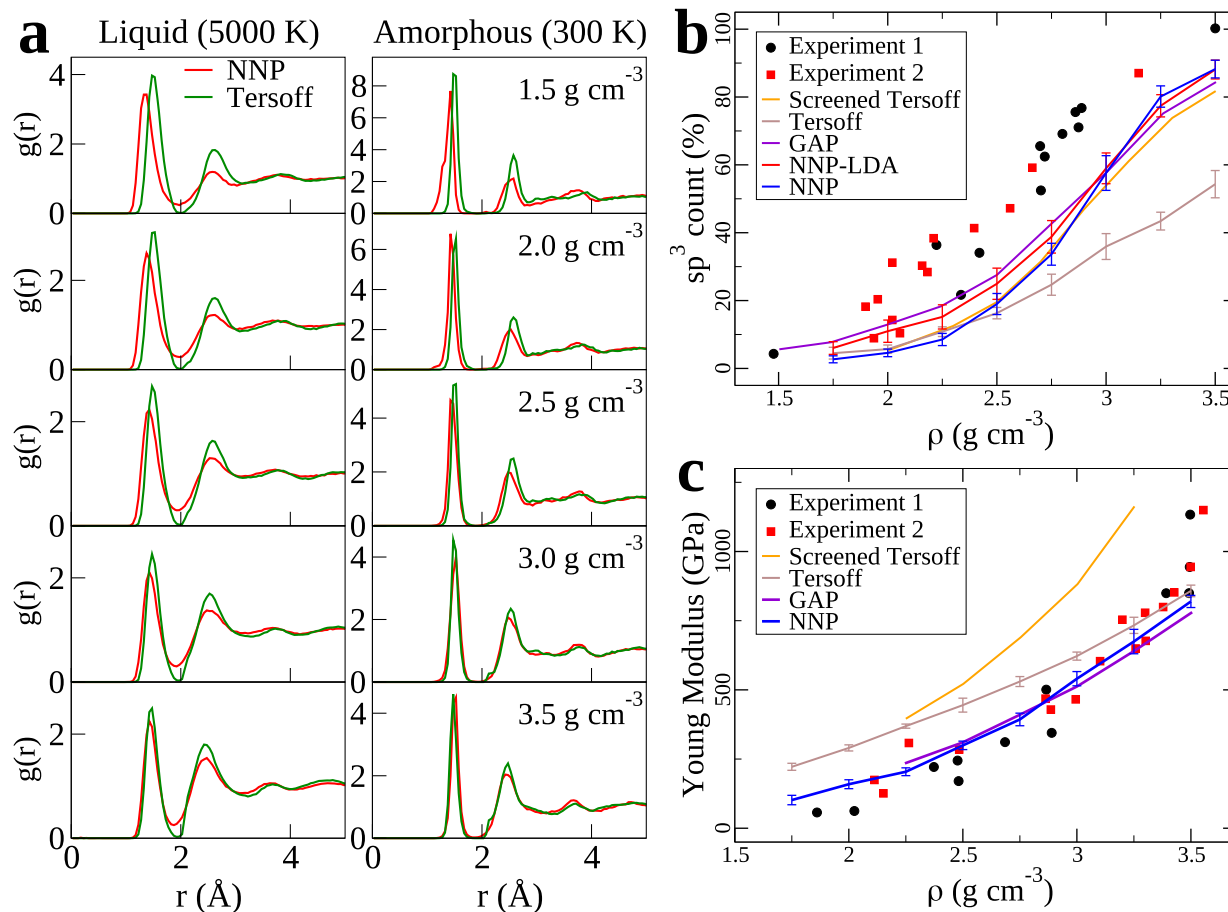


Fig. 6 Performance of the NNP on amorphous phases of carbon. **a** Radial distribution function for liquid (left) and amorphous (right) carbon, for our NNP and Tersoff potential, at increasing densities (top to bottom). **b** Percentage of tetrahedrally coordinated atoms in amorphous carbon structures as a function of density, comparing NNP with rVV10 and LDA-level, and Tersoff potential to results taken from ref. ²¹ for GAP and screened Tersoff potentials, as well as experimental results from refs. ^{42,43}. **c** Young modulus of amorphous carbon as a function of density for NNP at rVV10 level and Tersoff, compared to results taken from ref. ²¹ for GAP and screened Tersoff potentials, as well as experimental results from refs. ^{76,77}. Error bars represent standard deviation over ten random initialization of particle velocities of the melt-quench cycles at each density.

input space, to as high as an RMSE of 473 meV/atom. This is a confirmation of the common observation that the prediction error of a neural network is strongly dependent on the similarity of training and test environments⁴⁷. On the other hand, when the model is trained and tested using the complete set, a prediction RMSE of 22.1 meV/atom is obtained for energy, while, for the configurations within $D=0.05$ from diamond, the prediction RMSE is still considerably small, 7.7 meV/atom. The analysis for forces follows the same trend as energies. The RMSE values for energies and forces are given in the Supplementary Table I.

Hence, it can be deduced that, for a fixed network architecture, a trade-off must be struck between having small error on configurations similar to a reference structure, and obtaining reliable predictions for general configurations from the full potential energy surface. The other entries in these tables confirm this analysis: the more diverse the training set is, the more robust is the resulting potential outside its training basin. Therefore, for a reliable NNP for multiple C polymorphs, as the one targeted here, a diverse training set from a wide region of the potential energy surface is necessary.

In summary, in this work, we have presented a self-consistent technique for generating an accurate and transferable NNP. Since neural networks encode the physics of a system into their parametrization through data, the dataset plays a crucial role in the resulting NNP performance. The method described in this

work achieves a comprehensive dataset via balanced integration of evolutionary algorithm, unsupervised machine learning in the form of clustering, and MD. As the training dataset is central to all machine learning models, we believe this generation method may be adopted by and would be beneficial to other ML approaches as well.

The distance-based analysis also gives an a posteriori measure of the profound diversity of the final dataset achieved via the self-consistent method. MD together with EA and clustering successfully explores a wide range of configurations on equal footing so that the dataset shown in Fig. 7c covers energy and volume landscape rather homogeneously. This is in line with the observation that at each iteration dataset diversity increases and validation RMSE may also increase since the network is tasked with a more complex functional approximation problem.

The presented workflow requires minimum human intervention. As the potential is iteratively improved, even rough starting models could be utilized for the very first step, and we have shown that the converged potential does not carry the limitations of the initial model. Therefore, not only this workflow is ready for high-throughput automation schemes as envisioned in future of experimentation but it is also robust with respect to lack of previous information about a system, as is often the case with novel materials.

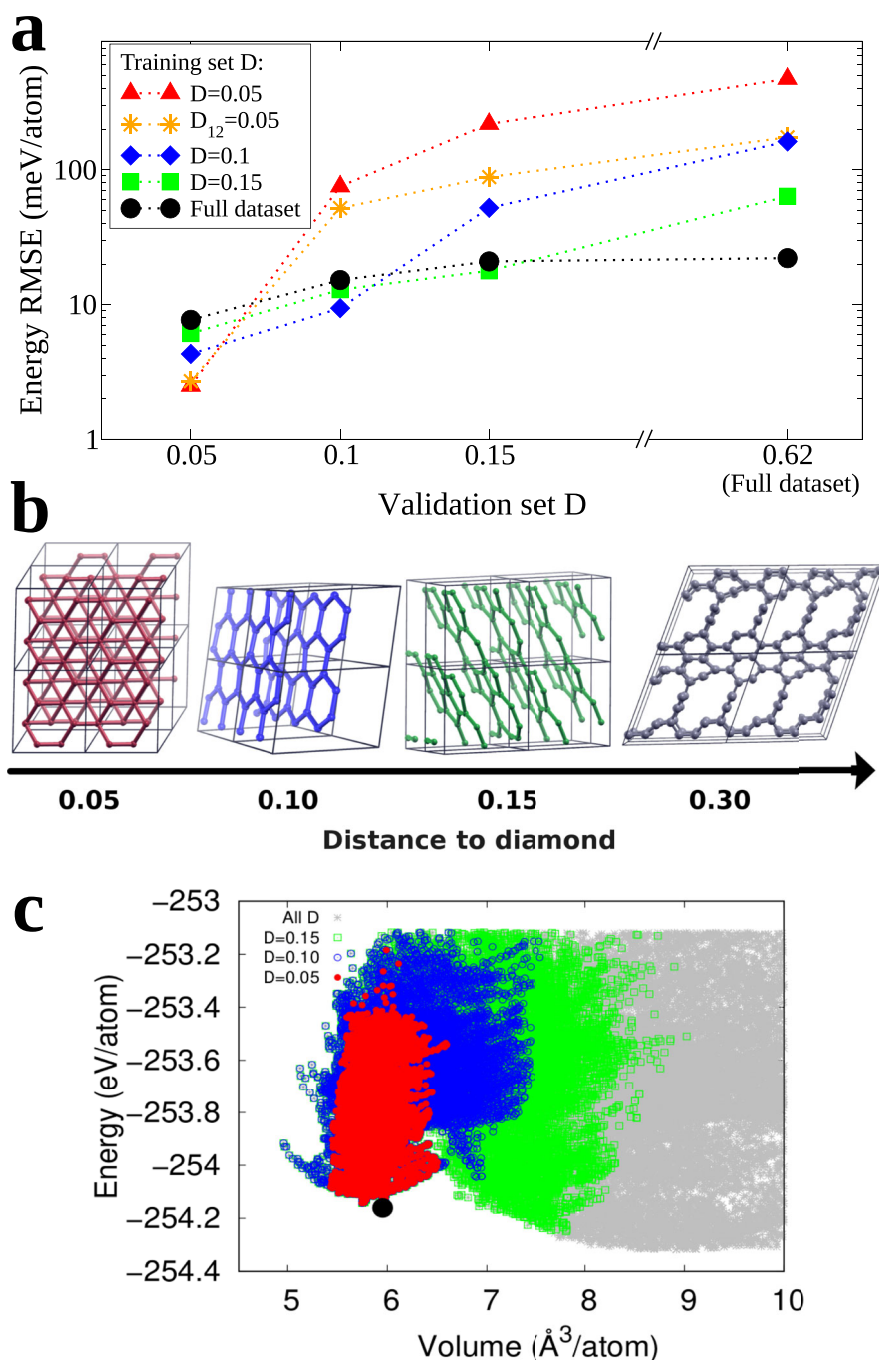


Fig. 7 The relationship between model transferability and the similarity of training and validation datasets. **a** Validation error of networks trained on different datasets as a function of the distance of the validation set from diamond. Numerical values are given in the Supplementary Table 1 for energies and forces. **b** Representative structures at given distances from diamond, the reference structure. The structures at 0.05 or lower are recognizably related to the reference, while at 0.10 and 0.15 compressed and/or defected layered structures are visible. At 0.30 and above, configurations with several double bonds and carbon chains appear. **c** Energy per atom as a function of volume for structures in the dataset, colored according to their distance cutoff D from diamond. The black dot corresponds to the reference diamond structure. The complete dataset includes structures with larger volume that are omitted here for clarity. The complete volume range is given in the Supplementary Fig. 10.

Many new materials with practical applications can be expected to be multicomponent systems. As the phase space of possible compounds grows larger and wildly unexplored, truly automated and unbiased approaches for an efficient exploration will become essential. We believe that our dataset generation approach (which can be coupled to any other ML approximator with multicomponent capability, e.g., ref. ⁴⁸) would be particularly suited to such systems. The workflow and the underlying neural network⁴⁹

and electronic structure codes are publicly available and are open-source.

The self-consistent NNP generation procedure is entirely system independent and we demonstrated its successful application to the challenging case of carbon for which classical and machine-learned potentials are abundant in literature. We show that for diamond, graphite, and graphene phases, NNP reported in this work performs considerably better than Tersoff, a classical

potential, and overall better than the existing machine-learned potentials for structural and elastic properties. Recently, a new GAP model trained on a large dataset with wide range of polymorphs was published⁵⁰. Based on our reproduction of ab initio reference and ML results of this model, a preliminary comparison is given in the Supplementary Fig. 11 and Supplementary Table II, and it is found that NNP performs as well as or better for all properties studied.

When predicting graphite phonon dispersion, NNP resulted in very good agreement for the majority of the modes, yet predicted instability for the very soft modes that relate to interlayer interaction. We have traced this behavior to the accuracy requirement in predicting such small forces. To increase accuracy using a fixed neural network architecture, we built the training set only with structures that are in the vicinity of graphite according to a fingerprint-based distance measure. The resulting potential provided accurate phonon frequencies but it showed poor generalization to a wider range of structures, compared to a more comprehensive potential trained on the entire dataset. This example highlights the need for a procedure to standardize the accuracy measure of NNPs and a more pressing need to build error estimate measures into the process of generating NNPs.

METHODS

Evolutionary algorithm for configuration space search

In iterative schemes, having a good starting point often means that a smaller number of iterations is needed to reach convergence. In a realistic use case scenario of NNPs, it is reasonable to expect that only a moderately well-fitting potential would be available as a starting point. To demonstrate this, we start the self-consistent cycle using a Li–C ReaxFF model to generate the initial configurations. This model is fit to DFT results with vdW correction and its details are set to describe well Li–C environments and defective graphite but not the wide range of solid C polymorphs considered in this work. We generate the initial configurations with 16 and 24 carbon atoms per unit cell at 0, 10, 20, 30, 40, and 50 GPa via EA as implemented in USPEX^{51,52}. At each pressure, we start with a population of 30 (50) randomly generated structures for the 16 (24) atoms per unit cell, and evolve it through the following evolutionary operations with the given ratios: heredity (two parent structures are combined) 50%, mutation (a distortion matrix is applied to a structure) 25%, or by generating new random structures 25%.

At each generation, structures are optimized in five successive steps: (a) constant pressure and temperature MD at 0.1 GPa and 50 K, respectively, for 0.3 ps with time step of 0.1 fs, (b) relaxation of cell parameters and internal coordinates until force components are $<0.26 \text{ eV } \text{Å}^{-1}$, (c) constant pressure and temperature MD at 0.1 GPa and 50 K, respectively, for 0.3 ps with time step of 0.1 fs, (d) relaxation of cell parameters and internal coordinates until force components are $<0.026 \text{ eV } \text{Å}^{-1}$, and (e) a final relaxation of cell parameters and internal coordinates until force components are $<0.0026 \text{ eV } \text{Å}^{-1}$.

Only the 70% most energetically stable parents were allowed to participate in the process of creating the new generation. In the heredity step, only sufficiently distinct structures (whose cosine distance, as defined in the next section, is greater than a given threshold) are considered as parents. This threshold is fixed at 0.008 in the first iteration, as it is small enough to allow deformed structures from the same polymorph to be parents. In order to enhance the diversity of the structures in the subsequent iterations, the threshold is increased to 0.05 so that the parents can be expected to be from different polymorphs.

Each structure search is evolved up to a maximum of 50 generations at the first iterations and 30 in the subsequent ones. The configuration space search performed this way produces a wide range of sp^2 , sp^3 and mixture of sp^2 and sp^3 structures, including defective layered structures.

Clustering

Initially, an unsupervised, bottom-up, distance-based hierarchical clustering approach with single linkage is used on all structures obtained with EA to identify the unique polymorphs. In the later iterations, clustering is applied only to those structures where NNP prediction differs from DFT ground-truth energy by more than 5 meV/atom. That way, polymorphs

that are already well described by NNP are not over-sampled. During clustering, to measure the similarity between structures, we use the fingerprint-based cosine distance defined in refs. ^{53,54}. In the case of a single species in the unit cell, and in its discretized form, the fingerprint of a configuration becomes:

$$F[k] = \frac{1}{2} \sum_{i \in \text{cell}} \sum_j \frac{\text{erf} \left[\frac{(k+1)\Delta - R_{ij}}{\sqrt{2}\sigma} \right] - \text{erf} \left[\frac{k\Delta - R_{ij}}{\sqrt{2}\sigma} \right]}{4\pi R_{ij}^2 \frac{N^2}{V} \Delta} - 1 \quad (2)$$

where the first sum runs over all atoms i in the unit cell and the second sum runs over all atoms j within a spherical cutoff radius R_{max} , and R_{ij} is the distance between atoms i and j . The numerator describes the integral of a Gaussian density of width σ over a bin of size Δ . N is the number of atoms in the unit cell and V is the unit cell volume.

The cosine distance between structures 1 and 2 is defined as:

$$D_{\text{cosine}}(1, 2) = \frac{1}{2} \left(1 - \frac{\mathbf{F}_1 \cdot \mathbf{F}_2}{|\mathbf{F}_1| |\mathbf{F}_2|} \right). \quad (3)$$

The dimension of the F -vector is set to $R_{\text{max}}/\Delta = 125$ with $R_{\text{max}} = 10 \text{ Å}$ and $\Delta = 0.08$ in this work. Two configurations closer to one another than a distance threshold are determined to belong to the same cluster. In this work the threshold is tuned to yield ~100–150 clusters at each step, which results in affordable computational cost for the remaining calculations of the self-consistent cycle.

Molecular dynamics (MD)

We manually select a representative structure from each cluster and perform a 0.5-ns classical NPT MD simulation with Nose–Hoover thermostat and barostat. In these simulations, the external conditions of pressure and temperature are ramped up from -50 GPa at 100 K , to 50 GPa at 1000 K in the course of 0.5 ns. The characteristic relaxation times of the thermostat and barostat are chosen as 50 and 100 fs, respectively. By sampling a snapshot of the dynamics every 5 ps, 100 configurations are selected. All MD simulations are performed with LAMMPS package⁵⁵. In addition, 440 randomly selected graphene atomic configurations from the libAtoms repository³⁶ are added to the selection. This set constitutes the set of structures where ab initio total energy calculations are then performed and added to the training set.

First principles calculations

The first principles calculations performed on all the structures visited during EA configuration space search and MD refinement described earlier employ the following parameters: plane wave basis set kinetic energy cutoff for wavefunctions and charge density is 80 and 480 Ry, respectively. The rVV10⁵⁶ exchange-correlation functional that incorporates non-local vdW correlations is employed. A Brillouin zone sampling with resolution of $0.034 \times 2\pi \text{ Å}^{-1}$ for the 3D carbon structures and $0.014 \times 2\pi \text{ Å}^{-1}$ for graphene is used. These parameters are found to yield 1 mRy/atom precision on diamond, graphite, and graphene. All DFT calculations were performed with the Quantum ESPRESSO package^{57,58}. Elastic properties are computed through the thermopw framework⁵⁹ while vibrational properties are obtained with PHON package⁶⁰.

In the first self-consistent iteration, the training set is made up of all generated structures lying within 10 eV from the lowest energy one. This results in a total of ~16,000 configurations. In the subsequent iterations of the self-consistent procedure, we use all configurations whose energy per atom is within 1.2 eV of the lowest one, these are added to the previously selected configurations, amounting to a total of about 30,000 configurations in the second and 60,000 configurations in the third and final iteration. From these configurations, 20% was set aside for validation and the remaining 80% was used in the NNP training.

Neural network architecture

In this work, we adopt the Behler–Parrinello approach to atomistic neural networks⁴⁴ where the total energy of a system of N atoms is defined as the sum of atomic energy contributions

$$E = \sum_{i=1}^N E_i(G_i), \quad (4)$$

where E_i is the energy contribution of an atom i , and G_i is its local environment descriptor vector. As described in detail in the next section, we choose descriptors with 144 components per atomic environment.

The contribution of an atom to the total energy is obtained by feeding its environment descriptor to the feed-forward all-to-all-connected neural network. Here we build a network with two hidden layers, with 64 and 32 nodes for the first and second layer, respectively, both with Gaussian activation function, and a single-node output layer with linear activation. The resulting network has a total of 11,393 parameters, i.e., $(144 \times 64) + (64 \times 32) + (32 \times 1) = 11,296$ weights and $64 + 32 + 1 = 97$ biases. The energy of each atom is then summed to obtain the total energy of the configuration. The force on each atom can be obtained analytically

$$\mathbf{F}_i = - \sum_j \sum_\mu \frac{\partial E_j}{\partial G_{j\mu}} \frac{\partial G_{j\mu}}{\partial \mathbf{R}_i} \quad (5)$$

where the atom index, j , runs over all the atoms within the cutoff distance of atom i , and index μ runs over the descriptor components.

During training, the weight and bias parameters W , are optimized with the Adam algorithm⁶¹ using gradients obtained by randomly selected subsets (minibatches) of data. The loss function of this stochastic optimization problem is defined as the sum of two contributions: one using the total energy value (Eq. (6)) and one using the force on each atom (Eq. (7)):

$$\mathcal{L}^E(W) = \sum_{c \in \text{batch}} (E_c^{\text{DFT}} - E_c(W))^2 + \exp \left[a \tanh \left(\frac{1}{a} \sum_{c \in \text{batch}} \left(\frac{E_c^{\text{DFT}} - E_c(W)}{N_c} \right)^2 \right) \right], \quad (6)$$

where E_c^{DFT} is the ground-truth total energy obtained via DFT and E_c is the NN prediction for total energy of a given configuration c , consisting of N_c atoms in the unit cell. The second part of this equation exponentially penalizes outliers while keeping the exponent normalized; a is a constant that allows to tune this penalty, $a = 5$ is used in this study. The force contribution to the loss is given by:

$$\mathcal{L}^F(W) = \gamma_F \sum_{c \in \text{batch}} \sum_{i=1}^{N_c} |\mathbf{F}_i^{\text{DFT}} - \mathbf{F}_i|^2, \quad (7)$$

where for any atom i of configuration c , $\mathbf{F}_i^{\text{DFT}}$ is the ground-truth force obtained via DFT, and \mathbf{F}_i is the NN prediction for it. γ_F is a user-defined parameter that controls the scale of this loss component. The results reported are obtained with γ_F equals 0.5. The relative error loss highlighted in "Results" is defined as

$$\mathcal{L}^F(W) = \gamma_F \sum_{c \in \text{batch}} \frac{1}{N_c} \sum_{i=1}^{N_c} \frac{|\mathbf{F}_i^{\text{DFT}} - \mathbf{F}_i|^2}{|\mathbf{F}_i^{\text{DFT}}|^2 + f_0^2}, \quad (8)$$

where f_0 is a regularizer constant, chosen as $f_0 = 260 \text{ meV } \text{\AA}^{-1}$ in this work.

An L_2 -norm regularization term is also added with a small coefficient $\gamma_R = 10^{-4}$ to prevent weights from becoming spuriously large

$$\mathcal{L}^R(W) = \gamma_R \frac{|W|^2}{2}. \quad (9)$$

The total loss is thus defined as:

$$\mathcal{L}(W) = \mathcal{L}^E(W) + \mathcal{L}^F(W) + \mathcal{L}^R(W). \quad (10)$$

All models are trained starting from random weights and a starting learning rate $\alpha_0 = 0.001$. The learning rate is decreased exponentially with optimization step t following the relationship $\alpha(t) = \alpha_0 e^{-t/\tau}$ with decay rate $r = 0.96$ and the decay step $\tau = 3200$. A batch size of 128 data points is used throughout the study.

Atomic environment descriptors

We use Behler–Parrinello symmetry functions⁴⁴ as local atomic descriptors. These functions include a two body and a three-body term, referred to as radial and angular descriptor, respectively. We use a modified version of the original angular descriptor⁴⁵ as implemented and detailed in PANNA package⁴⁶. The radial descriptor function is defined as:

$$G_i^{\text{Rad}}[s] = \sum_{j \neq i} e^{-\eta(R_{ij} - R_s)^2} f_c(R_{ij}), \quad (11)$$

where η and a set of Gaussian-centers R_s are user-defined parameters of the descriptor. The sum over j runs over all atoms whose distance R_{ij} from the central atom i is within the cutoff distance R_c . The cutoff function, f_c is

defined as:

$$f_c(R_{ij}) = \begin{cases} \frac{1}{2} \left[\cos\left(\frac{\pi R_{ij}}{R_c}\right) + 1 \right] & R_{ij} \leq R_c \\ 0 & R_{ij} > R_c. \end{cases} \quad (12)$$

The angular part of the descriptor with central atom i is defined as:

$$G_i^{\text{Ang}}[s] = 2^{1-\zeta} \sum_{j,k \neq i} (1 + \cos(\theta_{ijk} - \theta_s))^\zeta \times e^{-\eta(R_{ij}/2 + R_{ik}/2 - R_s)^2} \times f_c(R_{ij}) f_c(R_{ik}). \quad (13)$$

The sum runs over all pairs of neighbors of atom i , indexed as j and k , with distances R_{ij} and R_{ik} within the cutoff radius R_c , forming an angle θ_{ijk} with it. Here η , ζ , and the sets of θ_s and R_s are the user-defined parameters of the descriptor.

We note that the descriptor as written in Eq. (13) has discontinuous derivative with respect to atomic positions when atoms are collinear. To restore the continuity, we replace the $\cos(\theta_{ijk} - \theta_s)$ term with the following expression:

$$2 \frac{\cos(\theta_{ijk}) \cos(\theta_s) + \sqrt{1 - \cos(\theta_{ijk})^2 + \epsilon} \sin(\theta_s)^2 \sin(\theta_s)}{1 + \sqrt{1 + \epsilon} \sin(\theta_s)^2} \quad (14)$$

where we introduce a small normalization parameter, ϵ , such that the expression approaches $\cos(\theta_{ijk} - \theta_s)$ in the limit of $\epsilon \rightarrow 0$. In this work, $\epsilon = 0.001$ was used, while values between 0.001 and 0.01 were found to yield stable dynamics and equivalent network potentials for any practical purpose.

The radial descriptors are parametrized with $\eta = 16.0 \text{ \AA}^{-2}$, while 32 equidistant Gaussian centers, R_s , are distributed between 0.5 and 4.6 \AA . For the angular part $\eta = 10.0 \text{ \AA}^{-2}$, $\zeta = 23.0$, 8 equidistant R_s are distributed between 0.5 and 4.0 \AA and 14 θ_s are chosen between $\pi/28$ and $27\pi/28$ with spacing $\pi/14$. The cutoff R_c is 4.6 \AA for radial and 4.0 \AA for the angular descriptors, respectively. The resulting descriptor has a total of $32 + 14 \times 8 = 144$ components per atomic environment.

DATA AVAILABILITY

The neural network potential described in this work is released in PANNA⁴⁶ format compatible with several molecular dynamics packages via OPENKIM⁶². A native LAMMPS plugin version is also given in the Supplementary Material. The training and test data are available from the corresponding authors upon reasonable request.

Received: 24 May 2020; Accepted: 10 February 2021;

Published online: 14 April 2021

REFERENCES

- Hohenberg, P. & Kohn, W. Inhomogeneous electron gas. *Phys. Rev.* **136**, B864 (1964).
- Kohn, W. & Sham, L. J. Self-consistent equations including exchange and correlation effects. *Phys. Rev.* **140**, A1133–A1138 (1965).
- Chandrasekaran, A. et al. Solving the electronic structure problem with machine learning. *Nano Lett.* **5**, 22 (2019).
- Xie, T. & Grossman, J. C. Crystal graph convolutional neural networks for an accurate and interpretable prediction of material properties. *Phys. Rev. Lett.* **120**, 145301 (2018).
- Onat, B., Cubuk, E. D., Malone, B. D. & Kaxiras, E. Implanted neural network potentials: application to li-si alloys. *Phys. Rev. B* **97**, 094106 (2018).
- Kolsbjerg, E. L., Peterson, A. A. & Hammer, B. Neural-network-enhanced evolutionary algorithm applied to supported metal nanoparticles. *Phys. Rev. B* **97**, 195424 (2018).
- Cooper, A. M., Kästner, J., Urban, A. & Artrith, N. Efficient training of ann potentials by including atomic forces via taylor expansion and application to water and a transition-metal oxide. *npj Comput. Mater.* **6**, 54 (2020).
- Thompson, A., Swiler, L., Trott, C., Foiles, S. & Tucker, G. Spectral neighbor analysis method for automated generation of quantum-accurate interatomic potentials. *J. Comput. Phys.* **285**, 316–330 (2015).
- Zong, H., Pilia, 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).

10. Himanen, L., Geurts, A., Foster, A. S. & Rinke, P. Data-driven materials science: status, challenges, and perspectives. *Adv. Sci.* **6**, 1900808 (2019).
11. Nyshadham, C. et al. Machine-learned multi-system surrogate models for materials prediction. *npj Comput. Mater.* **5**, 51 (2019).
12. Kostichenko, 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).
13. 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).
14. 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).
15. Zubatyuk, R., Smith, J. S., Leszczynski, J. & Isayev, O. Accurate and transferable multitask prediction of chemical properties with an atoms-in-molecules neural network. *Sci. Adv.* **5**, eaav6490 (2019).
16. Tersoff, J. Empirical interatomic potential for carbon, with applications to amorphous carbon. *Phys. Rev. Lett.* **61**, 2879–2882 (1988).
17. van Duin, A. C. T., Dasgupta, S., Lorant, F. & Goddard, W. A. Reaxff: a reactive force field for hydrocarbons. *J. Phys. Chem. A* **105**, 9396–9409 (2001).
18. Rappe, A. K. & Goddard, W. A. Charge equilibration for molecular dynamics simulations. *J. Phys. Chem.* **95**, 3358–3363 (1991).
19. Khaliullin, R. Z., Eshet, H., Kühne, T. D., Behler, J. & Parrinello, M. Graphite-diamond phase coexistence study employing a neural-network mapping of the ab initio potential energy surface. *Phys. Rev. B* **81**, 100103 (2010).
20. Koukaras, E. N., Kalosakas, G., Galiotis, C. & Papagelis, K. Phonon properties of graphene derived from molecular dynamics simulations. *Sci. Rep.* **5**, 12923 (2015).
21. Deringer, V. L. & Csányi, G. Machine learning based interatomic potential for amorphous carbon. *Phys. Rev. B* **95**, 094203 (2017).
22. Wen, M. & Tadmor, E. B. Hybrid neural network potential for multilayer graphene. *Phys. Rev. B* **100**, 195419 (2019).
23. Artrith, N. & Behler, J. High-dimensional neural network potentials for metal surfaces: a prototype study for copper. *Phys. Rev. B* **85**, 045439 (2012).
24. Podryabinkin, E. V. & Shapeev, A. V. Active learning of linearly parametrized interatomic potentials. *Comput. Mater. Sci.* **140**, 171–180 (2017).
25. Artrith, N., Urban, A. & Ceder, G. Constructing first-principles phase diagrams of amorphous lixi using machine-learning-assisted sampling with an evolutionary algorithm. *J. Chem. Phys.* **148**, 241711 (2018).
26. Zhang, L., Lin, D.-Y., Wang, H., Car, R. & E. W. Active learning of uniformly accurate interatomic potentials for materials simulation. *Phys. Rev. Mater.* **3**, 023804 (2019).
27. Rowe, P., Csányi, G., Alfé, D. & Michaelides, A. Development of a machine learning potential for graphene. *Phys. Rev. B* **97**, 054303 (2018).
28. Bernstein, N., Csányi, G. & Deringer, V. L. De novo exploration and self-guided learning of potential-energy surfaces. *npj Comput. Mater.* **5**, 99 (2019).
29. Sivaraman, G. et al. Machine-learned interatomic potentials by active learning: amorphous and liquid hafnium dioxide. *npj Comput. Mater.* **6**, 104 (2020).
30. Shapeev, A. V. Moment tensor potentials: a class of systematically improvable interatomic potentials. *Multiscale Model Simul* **14**, 1153–1173 (2016).
31. Zuo, Y. et al. Performance and cost assessment of machine learning interatomic potentials. *J. Phys. Chem. A* **124**, 731–745 (2020).
32. Ma, Y. et al. Transparent dense sodium. *Nature* **458**, 182–185 (2009).
33. Bull, C. L. et al. ζ -Glycine: insight into the mechanism of a polymorphic phase transition. *IUCrJ* **4**, 569–574 (2017).
34. 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).
35. Deringer, V. L., Csányi, G. & Proserpio, D. M. Extracting crystal chemistry from amorphous carbon structures. *ChemPhysChem* **18**, 873–877 (2017).
36. Data repository for gaussian approximation potential. <http://www.libatoms.org/pub/Home/DataRepository>. (2018).
37. Perdew, J. P. & Wang, Y. Accurate and simple analytic representation of the electron-gas correlation energy. *Phys. Rev. B* **45**, 13244–13249 (1992).
38. Jacobson, P. & Stoupin, S. Thermal expansion coefficient of diamond in a wide temperature range. *Diam. Relat. Mater.* **97**, 107469 (2019).
39. Pozzo, M. et al. Thermal expansion of supported and freestanding graphene: lattice constant versus interatomic distance. *Phys. Rev. Lett.* **106**, 135501 (2011).
40. Evans, D. J. & Holian, B. L. The nose-hoover thermostat. *J. Chem. Phys.* **83**, 4069–4074 (1985).
41. Pastewka, L., Klemenz, A., Gumbusch, P. & Moseler, M. Screened empirical bond-order potentials for Si-C. *Phys. Rev. B* **87**, 205410 (2013).
42. Fallon, P. J. et al. Properties of filtered-ion-beam-deposited diamondlike carbon as a function of ion energy. *Phys. Rev. B* **48**, 4777–4782 (1993).
43. Schwan, J. et al. Tetrahedral amorphous carbon films prepared by magnetron sputtering and dc ion plating. *J. Appl. Phys.* **79**, 1416–1422 (1996).
44. Behler, J. & Parrinello, M. Generalized neural-network representation of high-dimensional potential-energy surfaces. *Phys. Rev. Lett.* **98**, 146401 (2007).
45. 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).
46. Lot, R., Pellegrini, F., Shaidu, Y. & Küçükbenli, E. Panna: Properties from artificial neural network architectures. *Comput. Phys. Commun.* **256**, 107402 (2020).
47. Bernstein, J., Vahdat, A., Yue, Y. & Liu, M.-Y. On the distance between two neural networks and the stability of learning, in *Advances in Neural Information Processing Systems*, eds: H. Larochelle and M. Ranzato and R. Hadsell and M. F. Balcan and H. Lin, 33, pp 21370–21381 (Curran Associates, Inc., 2020) <https://proceedings.neurips.cc/paper/2020/file/f4b31bee138ff5f7b84ce1575a738f95-Paper.pdf>.
48. Cusentino, M. A., Wood, M. A. & Thompson, A. P. Explicit multielement extension of the spectral neighbor analysis potential for chemically complex systems. *J. Phys. Chem. A* **124**, 5456–5464 (2020).
49. Panna: properties from artificial neural networks. <https://gitlab.com/PANNAdevs/panna>. (2020).
50. Rowe, P., Deringer, V. L., Gasparotto, P., Csányi, G. & Michaelides, A. An accurate and transferable machine learning potential for carbon. *J. Chem. Phys.* **153**, 034702 (2020).
51. Glass, C. W., Oganov, A. R. & Hansen, N. Uspex—evolutionary crystal structure prediction. *Comput. Phys. Commun.* **175**, 713–720 (2006).
52. Oganov, A. R. & Glass, C. W. Crystal structure prediction using ab initio evolutionary techniques: principles and applications. *J. Chem. Phys.* **124**, 244704 (2006).
53. Oganov, A. R. & Valle, M. How to quantify energy landscapes of solids. *J. Chem. Phys.* **130**, 104504 (2009).
54. Valle, M. & Oganov, A. R. Crystal fingerprint space—a novel paradigm for studying crystal-structure sets. *Acta Crystallogr. A* **66**, 507–517 (2010).
55. Plimpton, S. Fast parallel algorithms for short-range molecular dynamics. *J. Comput. Phys.* **117**, 1–19 (1995).
56. Sabatini, R., Gorni, T. & de Gironcoli, S. Nonlocal van der waals density functional made simple and efficient. *Phys. Rev. B* **87**, 041108 (2013).
57. Giannozzi, P. et al. Quantum espresso: a modular and open-source software project for quantum simulations of materials. *J. Phys. Condens. Matter* **21**, 395502 (2009).
58. Giannozzi, P. et al. Advanced capabilities for materials modelling with q quantum espresso. *J. Phys. Condens. Matter* **29**, 465901 (2017).
59. thermo_pw: ab-initio computation of material properties. https://dalcorso.github.io/thermo_pw/. (2020).
60. Alfé, D. Phon: a program to calculate phonons using the small displacement method. *Comput. Phys. Commun.* **180**, 2622–2633 (2009).
61. Kingma, D. P. & Ba, J. Adam: a method for stochastic optimization. <https://arxiv.org/abs/1412.6980> (2014).
62. Tadmor, E. B., Elliott, R. S., Sethna, J. P., Miller, R. E. & Becker, C. A. The potential of atomistic simulations and the knowledgebase of interatomic models. *JOM* **63**, 17 (2011).
63. Towns, J. et al. Xsede: accelerating scientific discovery. *Comput. Sci. Eng.* **16**, 62–74 (2014).
64. McSkimin, H. J. & Andreatch, P. Elastic moduli of diamond as a function of pressure and temperature. *J. Appl. Phys.* **43**, 2944–2948 (1972).
65. Zouboulis, E. S., Grimsditch, M., Ramdas, A. K. & Rodriguez, S. Temperature dependence of the elastic moduli of diamond: a Brillouin-scattering study. *Phys. Rev. B* **57**, 2889–2896 (1998).
66. Bosak, A., Krisch, M., Mohr, M., Maultzsch, J. & Thomsen, C. Elasticity of single-crystalline graphite: inelastic x-ray scattering study. *Phys. Rev. B* **75**, 153408 (2007).
67. Mohr, M. et al. Phonon dispersion of graphite by inelastic x-ray scattering. *Phys. Rev. B* **76**, 035439 (2007).
68. Seldin, E. J. & Nezbeda, C. W. Elastic constants and electron-microscope observations of neutron-irradiated compression-annealed pyrolytic and single-crystal graphite. *J. Appl. Phys.* **41**, 3389–3400 (1970).
69. Cooper, D. R. et al. Experimental review of graphene. *ISRN Condens. Matter Phys.* **2012**, 1–56 (2012).
70. Lee, C., Wei, X., Kysar, J. W. & Hone, J. Measurement of the elastic properties and intrinsic strength of monolayer graphene. *Science* **321**, 385–388 (2008).
71. Lee, J.-U., Yoon, D. & Cheong, H. Estimation of young's modulus of graphene by raman spectroscopy. *Nano Lett.* **12**, 4444–4448 (2012).
72. Curtarolo, S. et al. Aflowlib.org: a distributed materials properties repository from high-throughput ab initio calculations. *Comput. Mater. Sci.* **58**, 227–235 (2012).
73. de Pablo, J. J., Jones, B., Kovacs, C. L., Ozolins, V. & Ramirez, A. P. The materials genome initiative, the interplay of experiment, theory and computation. *Curr. Opin. Solid State Mater. Sci.* **18**, 99–117 (2014).
74. Draxl, C. & Scheffler, M. Nomad: the fair concept for big data-driven materials science. *MRS Bull.* **43**, 676–682 (2018).

75. Raju, M., Ganesh, P., Kent, P. R. C. & van Duin, A. C. T. Reactive force field study of li/c systems for electrical energy storage. *J. Chem. Theory Comput.* **11**, 2156–2166 (2015).
76. Schultrich, B., Scheibe, H.-J., Grandremy, G., Drescher, D. & Schneider, D. Elastic modulus as a measure of diamond likeness and hardness of amorphous carbon films. *Diam. Relat. Mater.* **5**, 914–918 (1996).
77. Schultrich, B., Scheibe, H.-J., Drescher, D. & Ziegele, H. Deposition of superhard amorphous carbon films by pulsed vacuum arc deposition. *Surf. Coat. Technol.* **98**, 1097–1101 (1998).

ACKNOWLEDGEMENTS

The work of E. Ka and E. Kü was supported by a DOE grant, BES Award DE-SC0019300. E. Kü, F.P., and S.d.G. are grateful for the financial support by European Union's Horizon 2020 research and innovation program under Grant agreement No. 676531 (project E-CAM). S.d.G. also acknowledges EU funding under Grant agreement No. 824143 (project MaX). This work used the high-performance computing resources of CINECA, SISSA, and FASRC Cannon cluster supported by the FAS Division of Science Research Computing Group at Harvard University. This work also used the Extreme Science and Engineering Discovery Environment (XSEDE), which is supported by National Science Foundation Grant number ACI-1548562⁶³, specifically it used Stampede2 at TACC through allocation TG-DMR120073.

AUTHOR CONTRIBUTIONS

E. Kü and S.d.G. designed and planned the study. E. Kü, F.P., and S.d.G. supervised all aspects of the project. R.L., F.P., Y.S., and E. Kü implemented the methodology into PANNA code, and performed extensive tests. Y.S. performed DFT calculations and constructed the ANN potentials. Y.S., F.P., S.d.G., and E. Kü analyzed the results. Y.S., E. Kü and S.d.G. led the manuscript writing. All authors contributed to discussions throughout the study and commented on the manuscript.

COMPETING INTERESTS

The authors declare no competing interests.

ADDITIONAL INFORMATION

Supplementary information The online version contains supplementary material available at <https://doi.org/10.1038/s41524-021-00508-6>.

Correspondence and requests for materials should be addressed to E.K. or S.d.G.

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) 2021