

ARTICLE OPEN



Performant implementation of the atomic cluster expansion (PACE) and application to copper and silicon

Yury Lysogorskiy¹, Cas van der Oord², Anton Bochkarev¹, Sarath Menon¹, Matteo Rinaldi¹, Thomas Hammerschmidt¹, Matous Mrovec¹, Aidan Thompson³, Gábor Csányi², Christoph Ortner⁴ and Ralf Drautz¹✉

The atomic cluster expansion is a general polynomial expansion of the atomic energy in multi-atom basis functions. Here we implement the atomic cluster expansion in the performant C++ code PACE that is suitable for use in large-scale atomistic simulations. We briefly review the atomic cluster expansion and give detailed expressions for energies and forces as well as efficient algorithms for their evaluation. We demonstrate that the atomic cluster expansion as implemented in PACE shifts a previously established Pareto front for machine learning interatomic potentials toward faster and more accurate calculations. Moreover, general purpose parameterizations are presented for copper and silicon and evaluated in detail. We show that the Cu and Si potentials significantly improve on the best available potentials for highly accurate large-scale atomistic simulations.

npj Computational Materials (2021)7:97; <https://doi.org/10.1038/s41524-021-00559-9>

INTRODUCTION

Atomistic modeling and simulation requires efficient computation of energies and forces. In recent years, machine learning (ML)-based interatomic potentials, parameterized to large data sets of reference electronic structure calculations, have provided particularly successful surrogate models of the atomic interaction energy. The ML models construct representations of atomic structure that are used in various regression algorithms to predict energies and forces.

The recently developed atomic cluster expansion (ACE)¹ provides a complete and efficient representation of atomic properties as a function of the local atomic environment in terms of many-body correlation functions. Because of the completeness of the ACE basis² these may be employed directly using linear regression for the computation of energies and forces. Furthermore, using simple nonlinear embedding functions, ACE can represent many classical as well as ML interatomic potentials. For example, the widely used family of embedded atom method (EAM)³ and Finnis–Sinclair (FS)⁴ potentials may be viewed as a lowest order ACE. Other properties, for example the moments of the density of states, may also be represented, and recursion or moments-based potentials like the bond-order potentials^{5,6} expanded in the form of an ACE.

Moreover, there are deep connections between ACE and several ML representations and formulations. The only other known complete parameterizations, the moment tensor potentials (MTP)⁷ and the ML potential of Seko et al.⁸, are both based on a body-ordered invariant polynomial basis and can be exactly represented by ACE by suitable choice of hyperparameters and an explicit linear transformation. In addition, the spectral neighbor analysis potential (SNAP)⁹, the atomic permutation-invariant potentials¹⁰, and descriptors such as the symmetry functions of Behler¹¹ and the smooth overlap of atomic positions (SOAP)¹² can be obtained as special cases or minor variations of the ACE formalism; see refs. ^{1,2,13} and Supplementary Methods for further details.

Here, we present the performant implementation of the atomic cluster expansion (PACE) enabling efficient evaluation of ACE

models within the LAMMPS molecular dynamics simulation software package¹⁴ (<https://www.lammps.org>). We demonstrate in Fig. 1 for two representative elements, Cu and Si, that PACE lowers the Pareto front of accuracy vs. computational cost that was established for several state-of-the-art ML potentials¹⁵. The details of how these were constructed are provided in the Supplementary Methods. While these benchmarks establish advanced computational performance, we also demonstrate the capacity of the ACE framework to develop highly accurate parameterizations: we present two parameterizations of interatomic potentials for Cu and Si that outperform available ML-based potentials in terms of performance, accuracy, and generalizability.

The fundamental building block from which ACE models are built are atomic properties φ_i which are expanded in terms of body-ordered functions from the set of neighbors of each atom i :

$$\varphi_i = \sum_{v=1}^{V_{\max}} \sum_{\mathbf{v}} \tilde{\mathbf{c}}_{\mathbf{v}} \sum_{j_1, \dots, j_v} \Phi_{\mathbf{v}}(\mathbf{r}_{j_1}, \dots, \mathbf{r}_{j_v}), \quad (1)$$

where $\Phi_{\mathbf{v}}$ are v -order basis functions (each involving coordinates of v neighbors), $\tilde{\mathbf{c}}_{\mathbf{v}}$ the model parameters, and \mathbf{v} the basis function indices. It appears as if this incurs an $O(N^v)$ computational cost, where N denotes the number of interacting neighbors; however, ACE exploits a much faster evaluation strategy that makes it possible to compute efficiently high body-order terms. This is achieved by (1) projecting the atomic density:

$$\rho_i(\mathbf{r}) = \sum_{j \neq i} \delta(\mathbf{r} - \mathbf{r}_{ji}), \quad (2)$$

on atomic basis functions, $\phi_{\mathbf{v}}(\mathbf{r})$, resulting in:

$$A_{iv} = \sum_{j \neq i} \phi_{\mathbf{v}}(\mathbf{r}_{ji}), \quad (3)$$

and (2) choosing a tensor product basis:

$$\Phi_{\mathbf{v}}(\mathbf{r}_1, \dots, \mathbf{r}_v) = \prod_{t=1}^v \phi_{v_t}(\mathbf{r}_{ti}), \quad (4)$$

¹ICAMS, Ruhr-Universität Bochum, Bochum, Germany. ²Engineering Laboratory, University of Cambridge, Cambridge, UK. ³Center for Computing Research, Sandia National Laboratories, Albuquerque, NM, USA. ⁴Department of Mathematics, University of British Columbia, Vancouver, BC, Canada. ✉email: ralf.drautz@rub.de

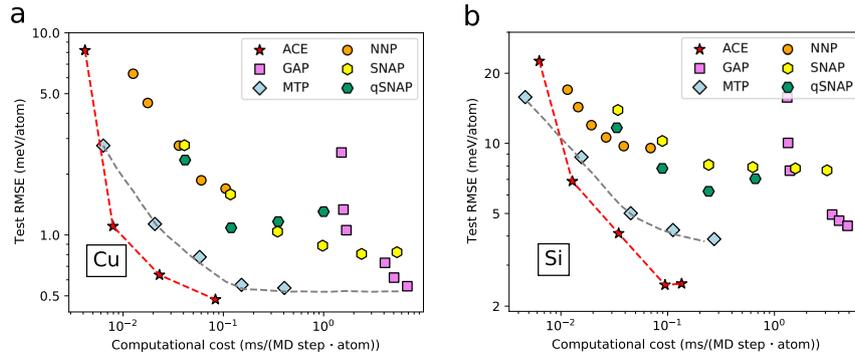


Fig. 1 ACE Pareto front. Test RMSE vs. computational cost for Cu (a) and Si (b) for ACE potentials compared to a recent benchmark study¹⁵. The timings from Zuo et al.¹⁵ were reduced by constant factors 0.55 (Cu) and 0.60 (Si) to correct for hardware differences and the ACE timings then overlaid.

with $\mathbf{v} = (v_1, v_2, \dots, v_\nu)$, which leads to¹:

$$\sum_{j_1, \dots, j_\nu} \Phi_{\mathbf{v}}(\mathbf{r}_{j_1, i}, \dots, \mathbf{r}_{j_\nu, i}) = \prod_{\ell=1}^{\nu} A_{i v_\ell}. \quad (5)$$

We call this reformulation the “density trick” (also used by Bartók et al.¹⁶ and Shapeev⁷ in formulating SOAP and MTP, respectively) and it results in the computational cost of an atomic property φ_i scaling linearly in N (due to evaluating the A_{ik}) and also linearly in ν (due to evaluating the correlations). Furthermore, we present a recursive evaluation scheme that avoids the ν -scaling altogether.

An ACE model may be defined in terms of several atomic properties $\varphi_i^{(p)}$, $p = 1, \dots, P$, for each atom i . For the simplest linear model of the potential energy one would use just one property, the atomic energy E_i :

$$E_i = \varphi_i^{(1)}. \quad (6)$$

A more elaborate model may generalize the pairwise repulsion and the pairwise density of the FS potential⁴ to arbitrary many-atom interactions:

$$E_i = \varphi_i^{(1)} - \sqrt{\varphi_i^{(2)}}. \quad (7)$$

In general, a large number of different atomic properties that are regarded as descriptors enter a nonlinear function:

$$E_i = \mathcal{F}(\varphi_i^{(1)}, \dots, \varphi_i^{(P)}), \quad (8)$$

where the nonlinearity \mathcal{F} may be explicit as in the FS model, or represent a general approximator such as artificial neural networks, as used by Behler and Parrinello¹⁷, or a kernel ridge regression model as used in the Gaussian approximation potential (GAP)¹⁶.

Different non-linearities \mathcal{F} may be used to incorporate physical or chemical insights in bond formation. Since the d -shell of copper is nearly full, angular contributions are generally small in the bulk. The unsaturated metallic bonds in the close-packed fcc ground state are modeled well by classical central-force functionals with nonlinear EAM or FS type embedding functions that effectively generate high body-order terms^{3,4,18}. Our parameterization for copper therefore starts from the FS representation of the energy, as in Eq. (7), but with the two atomic properties not limited to pairwise terms but including many-atom contributions that capture small angular contributions in the bulk and larger angular contributions in small clusters or two-dimensional structures.

On the other hand, the diamond structure of silicon is stabilized by angular contributions over close-packed structures, which highlights the importance of interactions beyond pairwise terms. In contrast to Cu the σ bonds in Si are nearly saturated¹⁹, which implies that to lowest order atomic energies in the open structures are linear in the number of bonds and a nonlinear embedding is

not required. Many different angle-dependent potentials have been developed for Si. Perhaps the best known are the Stillinger–Weber potential²⁰ with a linear three-body term and the Tersoff potential²¹ which includes nonlinear functions of three-body contributions. The most accurate potential for silicon to date, the SOAP-GAP model of Bartók et al.²² is an intrinsically high body-order potential. Here, we present a linear ACE for Si, which may be viewed as a generalization of this potential that includes all body-order interactions up to some maximum. In this way, ACE is employed in its basic form shown in Eq. (6), which simplifies the parameterization considerably and avoids implicit assumptions on the form of nonlinear terms that are often present in ML frameworks.

We carry out a detailed comparison of both our ACE parameterizations to the most reliable models available from literature. For Cu, we compare to the EAM potential by Mishin et al.²³, to a recent SNAP²⁴ parameterization as well as the GTINV⁸ ML potential. For Si, we compare to the GAP that was shown to reproduce a wide range of observable properties for crystalline, liquid, and amorphous Si phases²².

RESULTS AND DISCUSSION

Reference data

The parameterization for Cu was obtained by matching to the energies and forces of about 50,000 total energy calculations as obtained with density functional theory (DFT) using the PBE²⁵ functional as implemented in the FHI-aims code^{26,27}. The reference data included small clusters, bulk structures, surfaces and interfaces, point defects and their randomly modified variants. Part of the reference data has been briefly described in ref. ¹, but has been extended significantly for the present parameterization. For example, supercells for many elemental prototype structures with slightly displaced or missing atoms were added, as well as surfaces, interfaces and stacking faults (SF) with displaced atoms. Many structures were pulled apart until the atomic interactions vanished or compressed significantly. By far the most calculations were not relaxed to a force and stress free state. We employed `pyiron`²⁸ for generating part of the reference data.

The parameterization for Si was obtained by fitting to the same extensive silicon database GAP was fit to²². The database covers a wide range of configurations including crystalline structures, surfaces, vacancies, interstitials, and liquid phases. The DFT reference data were generated using the `CASTEP`²⁹ software package.

Parameterization and timing

We used different parameterization strategies for Cu and Si motivated by their different bond chemistry. In particular, the Si

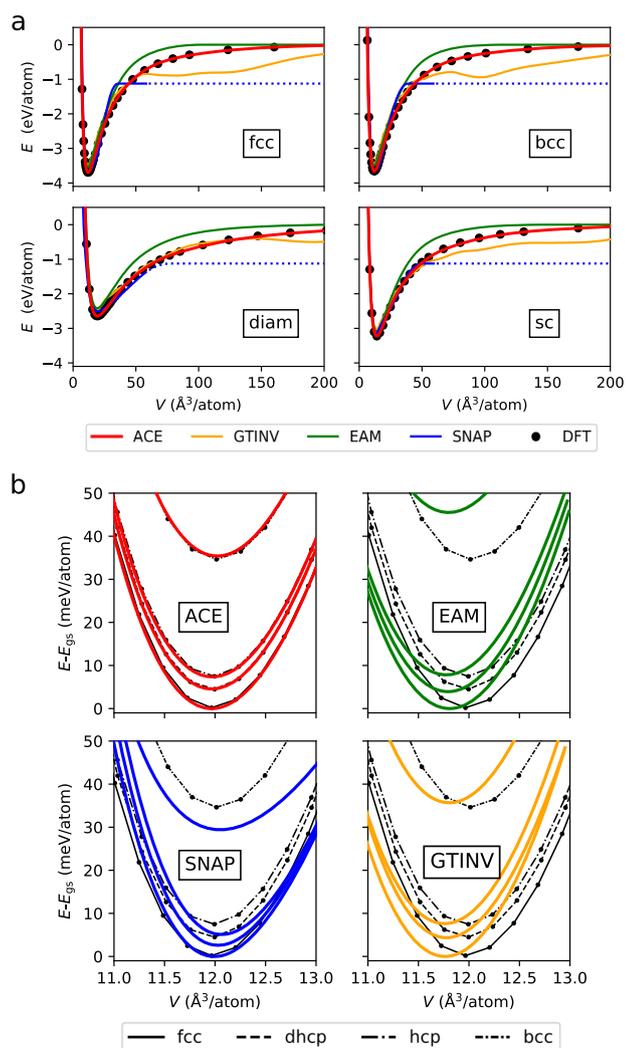


Fig. 2 Energy vs. volume for Cu. **a** Wide volume changes for fcc, bcc, diamond, and simple-cubic crystal structures. The SNAP cohesive energy was adjusted by a constant shift to match the fcc DFT data. **b** Low-energy bulk phases. Lines with symbols correspond to DFT reference data.

parameterization was obtained from solving a linear system of equations, whereas the Cu fit required nonlinear optimization. Energies and forces from the reference data were used in the parameterization. The Cu potential has a total number of 2072 parameters, of which 756 are expansion coefficients for each of the two densities, and 560 parameters are used for the radial functions. The DFT reference showed that interactions are smaller than 1 meV when atoms are further than the cutoff distance $r_c = 7.4 \text{ \AA}$ apart, when rigidly separating slabs. Parameter optimization led to a fit with an error of 3.2 meV/atom for structures that are within 1 eV of the ground state. This fit was then fine-tuned toward structures close to the ground state, which further decreased the error to 2.9 meV/atom and slightly increased the error of higher energy structures, with errors of 15 meV/ \AA per force component. For Si we used a total of 6827 basis functions parameterized as a linear model, with a maximum body order corresponding to $\nu = 4$. These basis functions were selected using the construction outlined in Dusson et al.² We show the silicon ACE matches the accuracy of the general purpose GAP potential introduced by Bartók et al.²² More specifically the energy error for the ACE model was found to be 1.81 meV/atom and errors of 75 meV/ \AA for each force component for structures within 1 eV

from the ground state. The corresponding errors for the GAP model are 1.25 meV/atom and 82 meV/ \AA on the silicon database presented in Bartók et al.²² For the Si ACE we used a cutoff distance of $r_c = 6.5 \text{ \AA}$ for the pair contribution and $r_c = 5.5 \text{ \AA}$ for the many-body part.

To evaluate the computational efficiency of PACE we carried out molecular dynamics (MD) simulations for face-centered cubic (fcc) Cu and diamond Si structures. We found that a single force call takes 0.32 and 0.80 ms/atom, respectively, for the Cu and Si ACE models (Timings were obtained on a single core of an Intel(R) Xeon(R) Gold 6132 CPU, using the GCC 7.3.0 compiler and LAMMPS version from 4 February 2020). These speeds are sufficiently fast for large-scale MD simulations and Monte Carlo sampling, for example, for the computation of phase diagrams. The efficiency of PACE is about two orders of magnitude slower than empirical potentials, cf. exemplary timings in μs per atom for Cu (EAM²³: 1.5, ADP³⁰: 7.0, MEAM³¹: 10) and Si (Tersoff³²: 2.6, MEAM³³: 18, ADP³⁴: 4.0).

Copper

The ACE for Cu has been comprehensively validated against DFT and available experimental data and compared to three other Cu potentials. The potentials we chose for the comparison were (1) the EAM potential of Mishin et al.²³, which exhibits an excellent overall accuracy and is considered as the reference Cu EAM potential, (2) the SNAP model of Li et al.²⁴, which was trained to strained crystalline as well as melted Cu phases obtained by ab initio MD, and (3) the ML interatomic potential Cu-gtinv-934 (GTINV) of Seko et al.⁸, which was fitted to an extended DFT database of 10^4 structures and reached RMSE values of 8.2 meV/atom. The EAM and SNAP potentials were computed through the OpenKIM API³⁵.

We evaluate the models for a broad range of structures and properties that not only exceed beyond the reference data but are also relevant for observable macroscopic behavior of Cu. Figure 2a gives an overview of the binding energy over large volume changes. ACE provides a very good match to the reference data at all distances, while the shorter range of EAM means that interatomic interactions are cut-off too early when the atoms are separated. The even shorter cutoff of SNAP leads to abrupt bond breaking, illustrating that the cohesive energy was not fitted in the construction of the potential and therefore one cannot apply the potential, for example, for gas phase condensation simulations. GTINV shows significant oscillations at larger interatomic distances.

A detailed analysis of structures that are energetically close to the fcc ground state is presented in Fig. 2b. All potentials reproduce correctly the structural order of fcc \rightarrow dhcp \rightarrow hcp \rightarrow bcc. The energy minima predicted by EAM and GTINV potentials are shifted to smaller volumes, which may be due to different DFT reference data. EAM and SNAP also show larger discrepancies for the fcc–bcc energy difference.

The elastic moduli for the ground state fcc structure are summarized in Table 1. ACE and EAM reproduce the DFT reference very well, while small deviations are observed for SNAP and slightly larger for GTINV. Similar outcomes are obtained for other bulk phases (see Supplementary tables) with ACE and EAM giving consistently the best agreement with DFT.

Despite not having fitted any phonon frequencies explicitly, ACE provides the best match to the reference DFT data. Supplementary Fig. 1 shows a comparison of phonon band structures and densities of states (DOS) for fcc Cu. The EAM and GTINV potentials overestimate the width of the DOS, while SNAP underestimates it. These conclusions apply also for the phonon DOS of other crystal structures that are shown in the Supplementary figures.

Table 1. Basic properties of Cu and Si.

Copper						
	ACE	EAM	GTINV	SNAP	DFT	Exp.
Elastic moduli (GPa)						
C_{11}	177	176	188	176	177	177 ⁵⁶
C_{12}	132	133	141	143	132	125 ⁵⁶
C_{44}	82	82	88	88	82	81 ⁵⁶
Surface energies γ_{surf} (J/m ²)						
(111)	1.36	1.24	7.81	1.29	1.36	–
(100)	1.51	1.35	9.50	1.48	1.51	–
(110)	1.59	1.48	7.82	1.56	1.57	–
Vacancy formation/migration, interstitial formation (eV)						
E_{vac}^f	1.12	1.28	1.13	1.37	1.07	1.27 ⁵⁷
E_{vac}^m	0.71	0.69	0.78	0.91	0.74	0.67–0.76 ⁵⁸
$I_{\text{db}(100)}$	3.14	3.12	3.25	2.80	3.10 ⁴⁰	2.8–4.2 ⁵⁹
I_{oct}	3.38	3.29	3.48	2.96	3.35 ⁴⁰	–
I_{tetr}	3.68	3.63	3.87	3.42	3.64 ⁴⁰	–
Stacking fault energies (mJ/m ²)						
γ_{ESF}	48	45	–679	29	43	–
γ_{ISF}	43	45	–1101	29	41	–
γ_{MAX}	826	771	1635	1139	826	–
γ_{MIDDLE}	500	479	510	562	494	–
γ_{TWIN}	22	22	–628	14	21	–
Melting temperature (K)						
T_{M}	1272	1324	–	1372	1251 ³⁷	1358 ⁶⁰
Silicon						
	ACE				GAP	DFT
Elastic moduli (GPa)						
B	80				83	82
C_{11}	142				145	147
C_{12}	50				52	50
C_{44}	70				69	73
Surface energies γ_{surf} (J/m ²)						
(111)	1.47				1.50	1.56
(100)	2.11				2.12	2.17
(110)	1.51				1.55	1.52
Vacancy/interstitial (eV)						
E_{vac}	3.72				3.73	3.67
I_{db}	3.55				3.59	3.66
I_{hex}	3.36				3.48	3.72
I_{tetr}	3.47				3.58	3.91
DFT reference obtained with FHI-aims ^{26,27} for Cu and CASTEP ²⁹ for Si.						

Transformations between different crystal structures present a sensitive test for any interatomic potential as both bond distances and bond angles are changed simultaneously. In addition, the associated changes in atomic coordination effectively scrutinize the screening of pairwise terms by many-atom contributions. As shown in Fig. 3, all potentials agree well with the reference DFT data for the tetragonal, trigonal, and hexagonal paths. However, only ACE provides an excellent quantitative interpolation for all structures along all considered transformation paths. Especially, the orthorhombic transformation, which can be regarded as a generalization of the Bain path³⁶, is challenging for the other potentials.

We used thermodynamic integration to evaluate the free energy of the solid and liquid phases of Cu. The free energies intersect at $T=1272$ K, about 20 K above the 1251 ± 15 K predicted by DFT³⁷. The EAM and SNAP melting temperature at $T=1325$ K and 1372 K are close to the experimental value of 1358 K. The prediction of the melting point with GTINV was not possible due to long evaluation times and the lack of a parallel implementation.

Figure 4a shows the thermal expansion as obtained from MD simulations in the NPT ensemble. All models agree well with the experimental data for temperatures up to 600 K and exhibit minor deviations at high temperatures.

Planar defects include internal interfaces, such as SF and grain boundaries (GBs), where the local atomic density does not vary significantly but bond angles change compared to bulk. In contrast, at free surfaces the bond angles remain mostly unaltered but the surface atoms lose about half of their neighbors. Typically, central-force models such as EAM provide a good description of structures and energies of GBs but cannot capture well the large local density changes at surfaces which usually leads to underestimation of surface energies.

The small energy differences between the close-packed fcc, hcp, and related structures in Cu imply small SF energies. ACE predicts the SF energies in very good agreement with DFT reference data, as shown in Table 1, with comparable predictions from EAM. GTINV predicts negative SF energies, hinting at a different ground state. SNAP provides SF energies with slightly larger deviations from the reference data.

The energies of several twin and twist symmetric GBs ($\Sigma = 3, 5, 9$) are compared in Fig. 4b to reference DFT data from the Materials Project database³⁸. As expected, all potentials predict the GB energies very accurately, which suggests good transferability of all models for environments with small local density variations.

As noted above, surfaces present a much more stringent test than GBs. ACE provides the best agreement with DFT reference data for all low-index surfaces, as shown in Fig. 4c, while both SNAP and EAM consistently underestimate the surface energies. For GTINV we observed a detachment of the top surface layers during relaxation which resulted in unphysically high surface energies that were excluded from the comparison.

In addition to the energetics of surfaces, we examined bond breaking in various atomic environments. Such tests have practical relevance as they are related to fracture, surface adsorption or vaporization. We designed three distinct decohesion tests that are schematically shown in Fig. 5. These tests compare bond dissociation in the Cu dimer, detachment of a Cu adatom from the (111) surface, and an ideal rigid decohesion of bulk Cu slabs that leads to the formation of two (111) free surfaces. As can be seen from Fig. 5, ACE is the only model that is able to describe quantitatively accurately the impact of the atomic environment on bond breaking. The presence of neighboring atoms leads to an effective screening of the interatomic bonds and their interaction ranges³⁹. The dimer and the adatom have no neighbors so that their interaction range is longer than the interaction between two surfaces whose atoms are surrounded by bulk. Given the simplicity of EAM, it provides a surprisingly good account of bond breaking in the very different environments, while GTINV and SNAP have problems with this test.

Properties of point defects, such as mono-vacancy and self-interstitial, are often included in the fitting dataset. Given that only unrelaxed vacancy configurations were part of the reference data, ACE reproduces the vacancy formation energy very well while the other potentials overestimate the DFT reference by 0.1–0.3 eV, see Table 1. The migration barrier is reproduced well, too, by the models, apart from SNAP that overestimates the barrier.

None of the interstitial configurations were included in the ACE training set, but ACE results are consistent with those of the other

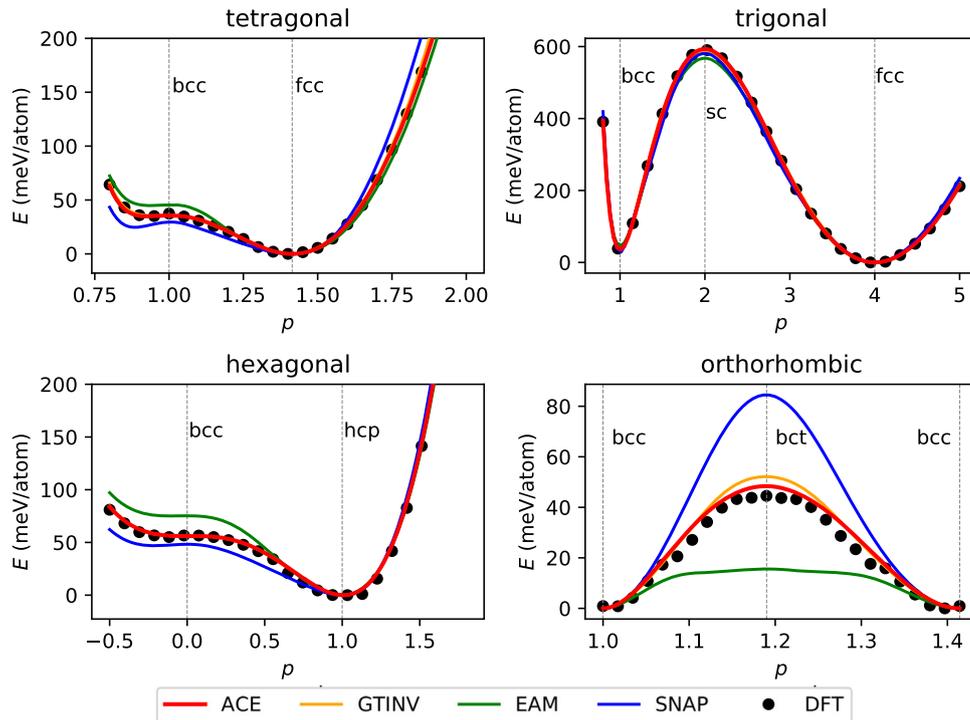


Fig. 3 Transformation paths for Cu. Shown are tetragonal, trigonal, hexagonal, and orthorhombic transformation paths.

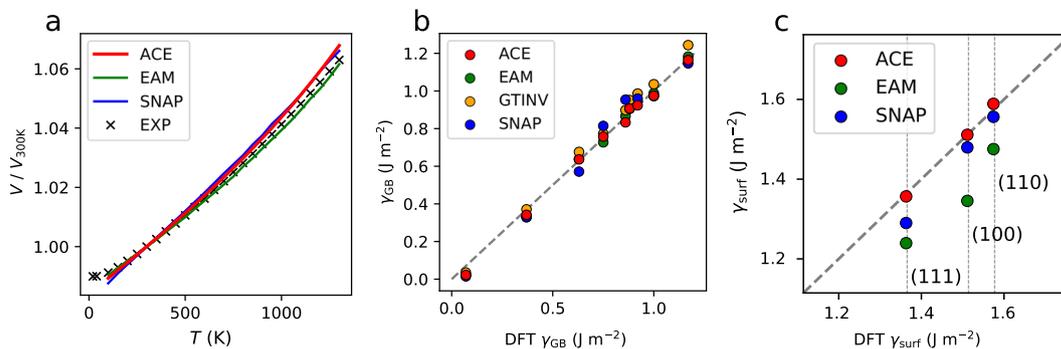


Fig. 4 Thermal expansion, grain boundary, and surface energies for Cu. **a** Thermal expansion for fcc, experimental data taken from ref. ⁶¹. **b** Twin and twist GBs, energies with respect to DFT reference from Zheng et al. ³⁸. **c** Surface energies.

potentials and together with EAM agree best with recent DFT results⁴⁰. The $\langle 100 \rangle$ dumbbell is predicted to have the lowest energy, followed by the octahedral and tetrahedral configurations. These predictions are consistent with those for other fcc metals⁴¹.

Small metallic clusters, important for catalysis and nanotechnology, usually form a large number of isomers with energies and structures often governed by subtle electronic structure contributions. For this reason, the predictions of the detailed energetics and structural stability is very challenging for interatomic potentials that are typically aimed at the description of bulk systems. We compared the predictions of ACE and the other models for three- and four-atomic clusters.

For the Cu trimer, the ground state structure is an isosceles triangle configuration while the linear trimer corresponds to an energy saddle point and is not dynamically stable⁴². In fact, the linear trimer transforms to a metastable configuration of a bent molecule with an obtuse angle of 130° . ACE is the only model that correctly reproduces the instability of the linear trimer and the existence of the metastable bent configuration. EAM predicts the equilateral triangle as the only stable configuration while for SNAP and GTINV both the linear trimer and the equilateral triangle are

stable configurations. The energy differences between the configurations are also reproduced most accurately by ACE, while the other models either significantly underestimate (GTINV) or overestimate (EAM, SNAP) the DFT values.

In the case of the tetramer, only ACE and GTINV give correctly the planar equilateral rhombus⁴² as the ground state, albeit GTINV shows also additional metastable configurations. Both EAM and SNAP favor incorrectly the close-packed tetrahedron which may originate from the lack of or weak angular contributions.

Planar 2D structures belong to a family of structures that is usually not included in the validation of interatomic potentials for bulk metals. It has been found recently that Cu is the only metal whose free-standing monolayers arranged in honeycomb, square, and hexagonal close-packed lattices are dynamically stable⁴³. We investigated in detail the 2D hcp lattice; both EAM and SNAP potentials show dynamic instabilities related to out-of-plane atomic displacements for the $3 \times 3 \times 1$ supercell that we used in our calculations (Supplementary Fig. 2). In contrast, DFT and ACE predict real phonon frequencies that confirm excellent transferability of ACE once more. We note that the 2D hcp structure could not be stabilized using the GTINV potential.

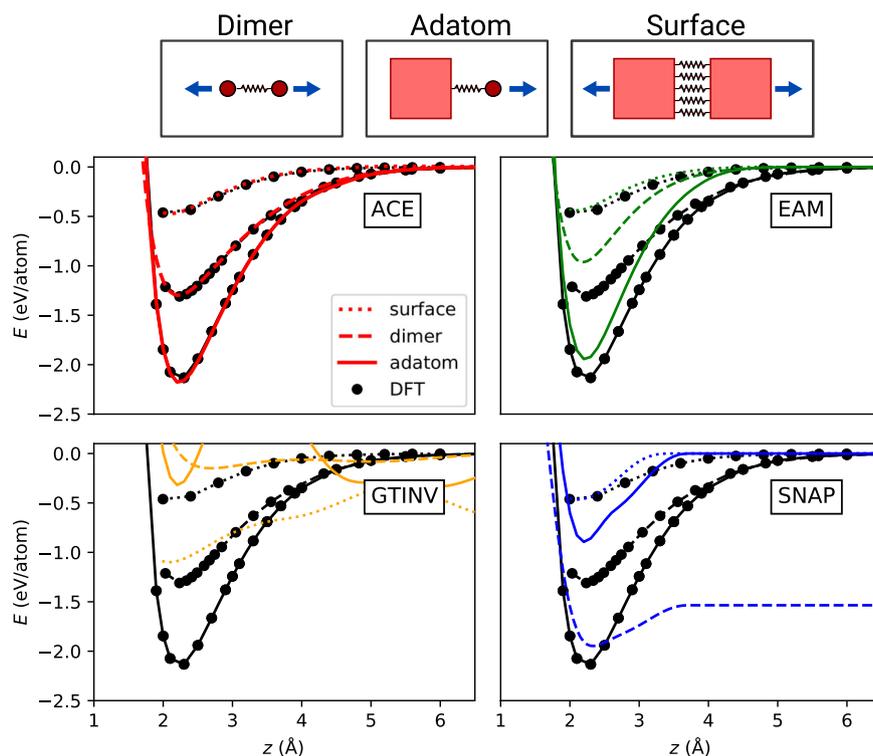


Fig. 5 Decoherence in different environments. Comparison of Cu dimer dissociation to a detachment of the “on-top” Cu adatom from the (111) surface and a rigid decohesion of two Cu bulk slabs along the $\langle 111 \rangle$ direction.

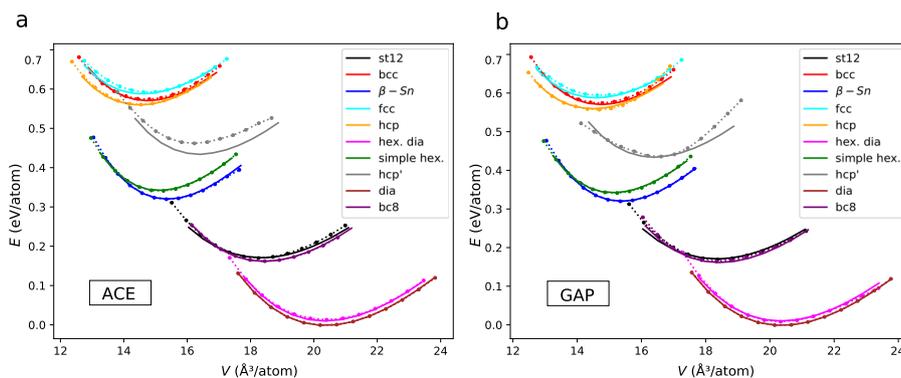


Fig. 6 Energy vs. volume for Si. Shown are bulk crystal lattices compared to DFT reference (solid). Both models were not explicitly fitted to the hcp' structure, requiring the ACE (a) and GAP (b) to extrapolate.

Silicon

The ACE for silicon was created by fitting to an extensive database first introduced to create a general purpose GAP model for silicon²². This GAP was shown to describe silicon accurately and to also be a qualitatively better interatomic potential than all other models tested, each best in their class: Stillinger–Weber^{20,44}, EDIP⁴⁵, Tersoff²¹, MEAM⁴⁶, DFTB⁴⁷, and ReaxFF⁴⁸. In this paper we show that the silicon ACE potential achieves the same accuracy as the GAP model, while being around 30 times faster in evaluation time and also better at extrapolating to unseen configurations.

The following presents a benchmarking of the ACE silicon potential on a wide range of properties including bulk, surface, liquid, and amorphous properties as well as a random structure search (RSS)⁴⁹ test.

The energies of the diamond, hexagonal diamond, β -Sn, bc8, st12, bcc, fcc, simple hexagonal, hcp, and hcp' are compared to DFT in Fig. 6. Excellent agreement with the DFT reference is

observed for all structures apart from hcp'. Si hcp has two minima²², the conventional hcp with $c/a \approx \sqrt{3}/2$, and hcp' with $c/a < 1$. The hcp' crystal structure is not contained in the DFT reference silicon database, however, both ACE and GAP predict the minimum. The GAP predicts the DFT reference energy at the minimum more accurately than ACE, while the latter gives a better estimate of the curvature.

The energy vs. volume curves for the silicon diamond and bcc are extended over a wide volume range in Fig. 7. Both potentials accurately describe the minima around 15 and 20 Å³/atom for bcc and diamond, as previously shown in Fig. 6. At larger volumes GAP exhibits unphysical high-energy minima. ACE does not show these minima and is close to the DFT reference, demonstrating better extrapolation compared to GAP. This extrapolative behavior is remarkable since there is no reference data at these large volumes as shown by the data density in the lower panel.

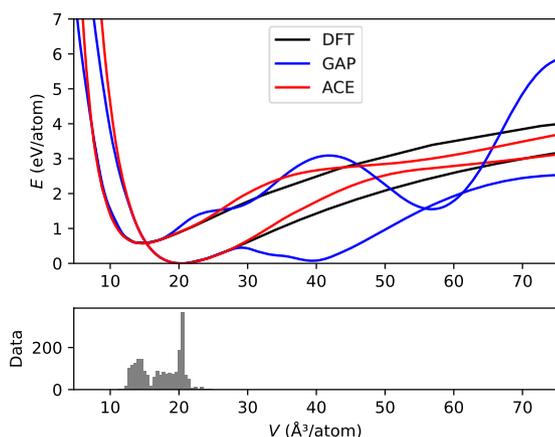


Fig. 7 Extrapolation for large volumes. Comparison of ACE and GAP to DFT for the energy-volume curves for diamond and bcc structures of Si. The lower panel shows the volume distribution of the reference data.

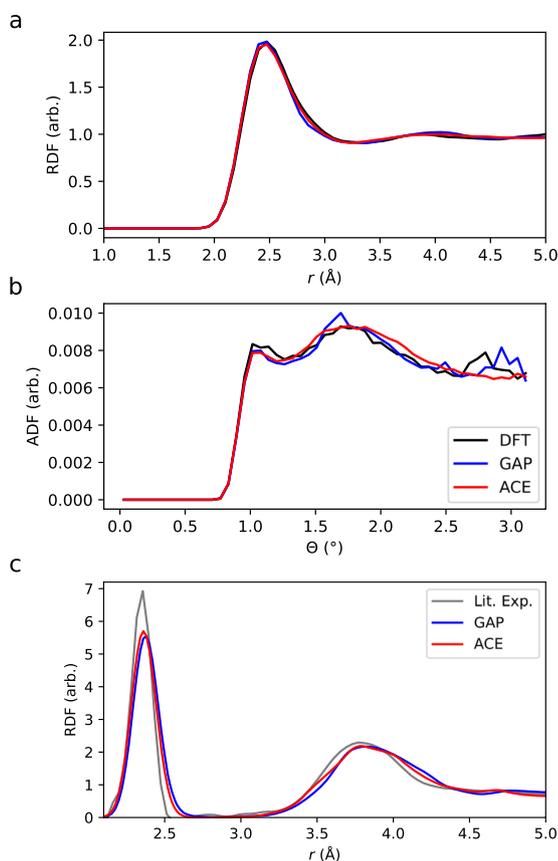


Fig. 8 Structural properties of liquid and amorphous Si. **a** Radial distribution function (RDF) and **b** angular distribution function (ADF) at $P = 0$ GPa and $T = 2000$ K. **c** Radial distribution function (RDF) for a 216-atom configuration generated by cooling liquid Si from 2000 to 500 K comparing ACE, GAP, and experimental data.

The elastic constants for Si in the diamond structure are summarized in Table 1. Both ACE and GAP match the DFT reference within a few percent.

Both ACE and GAP accurately describe the phonon spectrum in comparison to the DFT reference, with the band width of GAP showing a better match to DFT. The phonon dispersion of diamond for ACE and GAP and for other structures is shown in the Supplementary figures.

We investigated the thermal expansion, Grüneisen parameter, and heat capacity of ACE and GAP in the quasi-harmonic approximation as shown in Supplementary Fig. 4. Diamond silicon displays negative thermal expansion at low temperatures⁵⁰, which ACE models very well compared to the DFT reference, and more accurately than GAP. The heat capacity is described with almost perfect agreement, whereas the thermal expansion saturates at a slightly too high value for high temperature (for both GAP and ACE).

ACE was also tested in a liquid simulation on an eight-atom $2 \times 2 \times 2$ supercell (64 atom) and compared to GAP and the DFT reference. The radial distribution function (RDF) and angular distribution function were averaged over 20,000 MD steps (0.25 fs time step). The DFT reference data were generated using CASTEP averaging over 9700 MD steps (0.25 fs time step) taken at $T = 2000$ K, using a 200 eV plane-wave energy cutoff and 0.05 \AA^{-1} k-point spacing. The results are shown in Fig. 8 and demonstrate excellent agreement between ACE, GAP, and DFT reference.

Amorphous silicon is a tetrahedrally coordinated phase that forms upon rapid quenching from the melt. Here we quench a 216-atom sample of liquid Si from 2000 to 500 K at a rate of 10^{12} K/s with a 1 fs time step (1.5×10^6 steps) using the LAMMPS software. After the MD steps the final configuration was relaxed to a local minimum with respect to cell size and shape. The RDF of both GAP and ACE are compared to experimental results⁵¹ (since DFT results are not computationally feasible) in Fig. 8c. Both GAP and ACE accurately describe the first and second neighbor peaks, and no atoms in the range ($2.5 \text{ \AA} \leq r \leq 3.25 \text{ \AA}$).

The surface formation energy in the (100), (110), (111) directions are summarized in Table 1. ACE and GAP agree very well with the DFT reference.

Surface decohesion bridges two parts of the training database, from bulk crystal diamond to the unrelaxed (110) surface, see Supplementary Fig. 5. The unrelaxed surface and bulk crystal diamond were part of the database and accurately fitted, as well as some configurations along the path. The ACE energy is significantly smoother than GAP and closer to DFT reference.

The diamond vacancy formation energy and interstitial formation energies including tetragonal, hexagonal, and dumbbell are shown in Table 1. Both ACE and GAP predict point defect formation energies very well.

The lowest formation energy point defect is the “fourfold-coordinated defect” which consists of a bond rotation followed by a reconnection of all broken bonds⁵². We performed the following test using the ACE model: optimize the defect structure (using a 64 atom cell) with DFT, then re-optimize it with ACE, and finally compute the minimum energy transition path to the perfect crystal. When this test was performed with GAP in Bartók et al.²², no local minimum was found corresponding to the defect. With ACE however, we do find a local minimum, as shown in Fig. 9, where we also show the energy of the path evaluated with DFT and GAP. Remarkably, while both GAP and ACE make a similar error near the transition state, the ACE energy is significantly better for the relaxed defect structure, leading to the stabilization of the defect. Note that there are no configurations in the fitting database (which is identical for GAP and ACE) near the defect structure and the transition state, so again this shows the extrapolation power of the ACE model.

In the RSS^{49,53} test, randomized atomic configurations are relaxed, providing a view of the fitted potential energy surface for higher energy configurations. The RSS tests were performed on eight-atom configurations with close to cubic initial shapes and initial interatomic distances $>1.7 \text{ \AA}$. These configurations were then relaxed using the two-point steepest-descent method⁵⁴. The resulting energy per atom vs. volume per atom distribution is shown in Supplementary Fig. 6. ACE shows a similar distribution compared to DFT, with the diamond structure at the correct volume and a few structures up to 0.2 eV per atom higher at

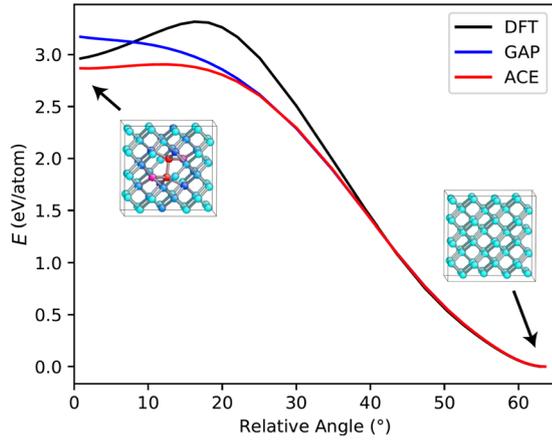


Fig. 9 Fourfold-coordinated defect. Minimum energy path connecting the fourfold defect to bulk crystal silicon evaluated using the ACE model in a 64 atom cell. We also show the energies of GAP and DFT on the exact same path.

comparable or somewhat larger volumes. A larger group of configurations is found at higher energies over a wider distribution of volumes. The density of states shows excellent agreement with DFT, as does the GAP model. This test is a strong discriminator between potentials, with all empirical potentials tested in Bartók et al.²² failing completely.

Discussion

We present a performant implementation of ACE in the form of the PACE code. We demonstrate that ACE, as implemented in PACE, shifts the Pareto front to higher accuracy and faster evaluation times, as compared to a number of ML potentials from ref.¹⁵. For our general purpose parameterizations of Cu and Si the CPU time per atomic force call is below 1 ms. As our implementation is fully compatible with LAMMPS, large scale simulations become possible, which we demonstrate through the computation of the free energies of liquid and solid phases for evaluating the melting temperature. PACE provides a simple interface for implementing nonlinear functions \mathcal{F} (Eq. (8)) as well as arbitrary radial functions which enables to adapt quickly to future ACE parameterizations.

We choose two distinct elements to illustrate parameterizations of ACE. Copper, for which classical potentials such as EAM are known to provide a good description of the interatomic interaction, and Si. For Si many different potentials were published to date and a recent GAP was shown to perform significantly better than other potentials²². We compare our Cu parameterization to a very good EAM potential and recent GTINV and SNAP potentials. The ACE for Si is compared to GAP.

For copper, EAM provides a very good description of the energy and ACE improves on this in particular for bonding environments that require angular contributions. Excellent extrapolation to new atomic environments is demonstrated, for example, for the phonons in a free-standing Cu monolayer. Furthermore, the longer cutoff of ACE enables us to reproduce bond breaking and making in accurate agreement with the DFT reference data. It appears that SNAP and GTINV were parameterized to selected reference data, which leads to deviations from DFT in several of our tests.

The Si ACE is comparable in accuracy to GAP, with a few key improvements. The ACE hypersurface is smoother than GAP, which is important in particular for extrapolation to large volumes as shown in Fig. 7. The improved smoothness can also be seen in the surface decohesion, showing behavior closely matching the DFT reference. Another example of the ACE extrapolation is the

fourfold defect which was highlighted in the original GAP paper, predicted erroneously to be unstable. However, ACE was fitted to the exact same DFT database and does predict a stable fourfold defect. Furthermore, it is notable that this Si ACE potential is ~30 times faster than the GAP of Bartók et al.²².

METHODS

We give detailed expressions for energies and forces and efficient algorithms for their evaluation in PACE in the following. For Cu and Si we employ distinct ACE forms, and different parameterization strategies follow for the two elements. The details of the parameterization strategy are provided in the Supplementary Methods.

Expressions for energy and forces

The energy of atom i is given by:

$$E_i = \mathcal{F}(\varphi_i^{(1)}, \dots, \varphi_i^{(p)}), \quad (9)$$

where \mathcal{F} is a general nonlinear function that may be supplied. Each atomic property $\varphi_i^{(p)}$ is given by an ACE expansion, which is obtained as follows: given the relative neighbor positions $\mathbf{r}_{ji} = \mathbf{r}_j - \mathbf{r}_i$, $r_{ji} = |\mathbf{r}_{ji}|$ and directions $\hat{\mathbf{r}}_{ji} = \mathbf{r}_{ji}/r_{ji}$, we first evaluate the atomic base:

$$A_{i\mu nlm} = \sum_j \delta_{\mu\mu_j} \phi_{\mu\mu_j nlm}(\mathbf{r}_{ji}), \quad (10)$$

where the one-particle basis ϕ is given in terms of spherical harmonics $Y_{lm}(\hat{\mathbf{r}}_{ji})$ and radial functions $R_{nl}^{\mu\mu_j}(r_{ji})$ by:

$$\phi_{\mu\mu_j nlm} = R_{nl}^{\mu\mu_j}(r_{ji}) Y_{lm}(\hat{\mathbf{r}}_{ji}). \quad (11)$$

Permutation-invariant many-body basis functions are obtained by forming products:

$$\mathbf{A}_{i\mu nlm} = \prod_{t=1}^v A_{i\mu_t n_t l_t m_t}. \quad (12)$$

The body order of the products is $v+1$ and the species of atom i is μ_i . The vectors $\boldsymbol{\mu}$, \mathbf{n} , \mathbf{l} , and \mathbf{m} have length v and contain atomic species, radial function indices, and spherical harmonics indices, respectively. The ACE expansion of an atomic property $\varphi_i^{(p)}$ is now given by:

$$\varphi_i^{(p)} = \sum_{\mu nlm} \tilde{c}_{\mu, \mu nlm}^{(p)} \mathbf{A}_{i\mu nlm}, \quad (13)$$

with expansion coefficients $\tilde{c}_{\mu, \mu nlm}^{(p)}$ and lexicographically ordered indices μnlm .

The coefficients $\tilde{c}_{\mu, \mu nlm}^{(p)}$ are not free model parameters to be fitted since the \mathbf{A} basis does not satisfy all required symmetries. An isometry invariant basis \mathbf{B} is obtained by coupling elements of the \mathbf{A} basis through the generalized Clebsch–Gordan coefficients, $\mathbf{B} = \mathbf{C}\mathbf{A}$, which yields a linear model:

$$\varphi_i = \mathbf{c}^T \mathbf{B} = \mathbf{c}^T \mathbf{C}\mathbf{A} \equiv \tilde{\mathbf{c}}^T \mathbf{A}, \quad (14)$$

from which we obtain $\tilde{\mathbf{c}} = \mathbf{C}^T \mathbf{c}$. The \mathbf{c} coefficients are the free model parameters that are optimized in the fit. We refer to Drautz¹, Dusson et al.², Drautz¹³ for details. It is helpful to think of the expansion coefficients $\tilde{c}_{\mu, \mu nlm}^{(p)}$ as satisfying linear constraints that ensure invariance of the properties φ_i and hence of the energy under rotation and inversion.

The force on atom k is written as:

$$\mathbf{F}_k = \sum_i (\mathbf{f}_{ik} - \mathbf{f}_{ki}), \quad (15)$$

and the pairwise forces \mathbf{f}_{ki} obtained using an adjoint method^{1,13}:

$$\begin{aligned} \mathbf{f}_{ki} &:= \nabla_{\mathbf{r}_i} E_i \\ &= \sum_{\mu nlm} \omega_{\mu nlm} \nabla_{\mathbf{r}_i} A_{i\mu nlm} \end{aligned} \quad (16)$$

$$= \sum_{nlm} \omega_{i\mu_k nlm} \nabla_k \phi_{\mu_k \mu_i nlm}(\mathbf{r}_{ki}), \quad (17)$$

where the adjoints $\omega_{\mu n l m}$ are given by:

$$\omega_{\mu n l m} = \sum_{\mu n l m} \Theta_{\mu n l m} \sum_t d\mathbf{A}_{\mu n l m}^{(t)}, \quad (18)$$

$$d\mathbf{A}_{\mu n l m}^{(t)} = \delta_{\mu\mu_t} \delta_{nn_t} \delta_{ll_t} \delta_{mm_t} \prod_{s \neq t} A_{\mu_s n_s l_s m_s}, \quad (19)$$

$$\Theta_{\mu n l m} = \sum_p \frac{\partial \mathcal{F}}{\partial \varphi_i^{(p)}} \tilde{\mathbf{c}}_{\mu, \mu n l m}^{(p)}. \quad (20)$$

A straightforward opportunity for optimization arises due to the fact that the product basis functions fulfill:

$$\text{Re}(\mathbf{A}_{\mu n l m}) = (-1)^{\sum_t m_t} \text{Re}(\mathbf{A}_{\mu n l -m}), \quad (21)$$

and $\sum_t m_t = 0$ for rotational invariance. As we are interested in a real-valued expansion, this identity is exploited by combining the $\mathbf{A}_{\mu n l m}$ and $\mathbf{A}_{\mu n l -m}$ and thus reducing the computational effort for evaluating the product basis by nearly 50%.

Similarly, when evaluating the forces only the real part needs to be evaluated, as the imaginary part has to add up to zero. Since:

$$\nabla \phi_{\mu, \mu, n l -m}(\mathbf{r}_{ji}) = (-1)^m (\nabla \phi_{\mu, \mu, n l m}(\mathbf{r}_{ji}))^*, \quad (22)$$

and therefore:

$$\omega_{\mu, \mu, n l -m} = (-1)^m (\omega_{\mu, \mu, n l m})^*, \quad (23)$$

one can limit the force evaluation to:

$$\begin{aligned} \mathbf{f}_{ki} &= \sum_{n l, m=0} \text{Re}(\omega_{\mu, \mu, n l 0}) \text{Re}(\nabla_k \phi_{\mu, \mu, l 0}(\mathbf{r}_{ki})) \\ &+ 2 \sum_{n l, m>0} \text{Re}(\omega_{\mu, \mu, n l m}) \nabla_k \phi_{\mu, \mu, l m}(\mathbf{r}_{ki}), \end{aligned} \quad (24)$$

which saves about 75% of the multiplications compared to fully evaluating all complex terms.

Algorithms

A PACE model is specified through four ingredients:

- (1) specification of the radial basis, typically as splines or through a polynomial recursion
- (2) a list of basis functions identified through $\mu n l m$ for each required order v
- (3) the corresponding expansion coefficients $\tilde{\mathbf{c}}_{\mu, \mu n l m}^{(p)}$
- (4) The nonlinearity $\mathcal{F}(\varphi_i^{(1)}, \dots, \varphi_i^{(p)})$ and its derivatives $\partial \mathcal{F} / \partial \varphi_i^{(p)}$

To formulate the evaluation algorithms it is convenient to reorganize the basis specification into a “compressed” format. First, we enumerate the list of one-particle basis functions and the atomic base A by identifying:

$$\mathbf{v} \equiv (\mu, n, l, m), \quad \text{and } A_{i\mathbf{v}} \equiv A_{\mu n l m}. \quad (25)$$

A tuple $\mathbf{v} = (v_1, \dots, v_v)$ can then be identified with $\mu n l m$ and specifies a corresponding many-body basis function:

$$\mathbf{A}_{i\mathbf{v}} = \prod_{\alpha=1}^r A_{i v_\alpha} \equiv \mathbf{A}_{\mu n l m}. \quad (26)$$

An atomic property φ_i can now be written more succinctly as:

$$\varphi_i^{(p)} = \sum_{\mathbf{v}} \tilde{\mathbf{c}}_{\mu, \mathbf{v}}^{(p)} \mathbf{A}_{i\mathbf{v}}. \quad (27)$$

This format condenses the notation as well as simplifies the basis specification, now given by (1) a list of one-particle basis functions indexed by \mathbf{v}_i and (2) a list of many-body basis functions, each represented by a tuple $\mathbf{v} = (v_1, \dots, v_v)$ where the length v specifies the interaction order.

The energy and force for a given atom i are obtained in five steps:

- (1) Evaluate atomic base $A_{\mu n l m}$, Eq. (10).
- (2) Evaluate product basis $\mathbf{A}_{i\mathbf{v}}$, Eq. (12) and properties $\varphi_i^{(p)}$, Eq. (13).
- (3) Obtain energy E_i , Eq. (9), and its derivatives with respect to the properties $\varphi_i^{(p)}$.
- (4) Compute product basis function derivatives $d\mathbf{A}_{i\mathbf{v}}^{(t)}$, Eq. (19), and adjoints $\omega_{\mu n l m}$, Eq. (18).
- (5) Assemble forces \mathbf{f}_{ji} , Eq. (17).

In the following we summarize algorithms for an efficient implementation.

Although we will not go into details of performance-oriented code optimizations, we briefly mention the four most important ingredients: (1) recursive algorithms to evaluate the polynomial, radial and spherical basis sets; (2) contiguous memory layout for the many-body basis specification; (3) recursive evaluation of the many-body basis; and (4) reducing the basis size and force evaluation by exploiting that the expansions are real; cf. Eqs. (21) and (24).

First the radial functions, spherical harmonics and their respective gradients are obtained. The spherical harmonics are computed in Cartesian coordinates directly¹³. Then the atomic base is evaluated. For $-l \leq m \leq 0$, we exploit:

$$A_{\mu n l -m} = (-1)^m A_{\mu n l m}^*. \quad (28)$$

Note that only for evaluations of the atomic base we need to work in $\mu n l m$ notation.

Algorithm 1. Atomic base A

```

Aμnlm = 0
for j ← neighbors of atom i do
  μ = type of atom j
  compute Rnl, dRnl, Ylm, dYlm
  for n, l, (m ≥ 0) do
    Aμnlm += Rnl · Ylm
  end for
end for
for n, l, (m > 0) do
  Aμnl-m = (-1)m Aμnlm*
end for

```

For numerical efficiency all pairwise radial functions can be represented as splines with several thousands interpolation points, which makes it possible to implement arbitrary radial basis functions efficiently.

Next the product basis functions \mathbf{A} and their derivatives Eq. (19) are set up.

Algorithm 2. Product basis functions \mathbf{A}

```

1: φi(p) = 0
2: for v ← basis, v = len(v) do
3:   Ai v = Re(∏t=1v Ai vt)
4:   for p do
5:     φi(p) += c̃μ, v(p) · Ai v
6:   end for
7: end for

```

The atomic properties $\varphi_i^{(p)}$ are computed following Eq. (13). Because $\mathbf{A}_{i\mathbf{v}}$ is used only to construct the $\varphi_i^{(p)}$, it need not be stored. Next, the energy $E_i = \mathcal{F}(\varphi_i^{(1)}, \dots, \varphi_i^{(p)})$ and its derivatives $\partial \mathcal{F} / \partial \varphi_i^{(p)}$ are obtained.

Once the derivatives $\partial \mathcal{F} / \partial \varphi_i^{(p)}$ are known, the adjoints $\omega_{\mu n l m}$ are computed following Eq. (18).

Algorithm 3. Adjoints ω

```

ωik = 0
for v ← basis, v = len(v) do
  Θi v = ∑p ∂F / ∂φi(p) · c̃μ, v(p)
  compute dAi v t, t = 1, ..., v (Alg. 4)
  for t ← 1, ..., v do
    ωi v t += Θi v · Re(dAi v t)
  end for
end for

```

Here, Θ_{iV} and $d\mathbf{A}_{iV}$ are only required locally and stored in a temporary variable. The derivatives $d\mathbf{A}_{iV}$ can be computed via backward differentiation with cost that scales linearly in v instead of the $O(v^2)$ scaling for a naive implementation. This important optimization is implemented as follows.

Algorithm 4. Compute $d\mathbf{A}_{iV}$, $t = 1, \dots, v$

```

1:  $\mathbf{A}^{\text{fwd}} = 1$ ;  $d\mathbf{A}_{iV1} = 1$ 
2: for  $t \leftarrow 1, 2, \dots, v - 1$  do
3:    $\mathbf{A}^{\text{fwd}*} = A_{iV_t}$ 
4:    $d\mathbf{A}_{iV(t+1)} = \mathbf{A}^{\text{fwd}}$ 
5: end for
6:  $\mathbf{A}^{\text{bwd}} = 1$ 
7: for  $t \leftarrow v, v - 1, \dots, 2$  do
8:    $\mathbf{A}^{\text{bwd}*} = A_{iV_t}$ 
9:    $d\mathbf{A}_{iV(t-1)}^* = \mathbf{A}^{\text{bwd}}$ 
10: end for

```

The computation of $d\mathbf{A}$ can be slightly improved by removing multiplications by one inside the loop. With that optimization the number of multiplications scales as $3v - 5$ for $v \geq 2$.

The gradients may now be obtained from Eq. (24).

Algorithm 5. Compute \mathbf{f}_{ki}

```

1:  $\mathbf{f}_{ki} = 0$ 
2: for  $n, l, (m \geq 0)$  do
3:   for  $k \leftarrow$  neighbors of atom  $i$  do
4:      $\mathbf{f}_{ki} += (2 - \delta_{m0}) \text{Re}(\omega_{i\mu_k n l m} \nabla_k \phi_{\mu_k i n l m}(\mathbf{r}_{ki}))$ 
5:   end for
6: end for

```

The overall computational cost is composed of two essentially independent contributions: (1) the evaluation of the atomic base A requires $O(N \cdot \#A)$ evaluations; and (2) the evaluation of the correlations \mathbf{A} requires $O((v_{\max} + \#\varphi) \cdot \#\mathbf{A})$ evaluations. That is, the overall cost scales linearly in the number of neighbors N , the maximum correlation order v_{\max} and also linearly in the number of properties $\varphi_i^{(p)}$. We will see next that the v -dependence can be further reduced with an alternative evaluation scheme.

Recursive evaluator

In most cases, the evaluation of the product basis functions and their derivatives (Algorithms 2, 3, and 4) are the computational bottleneck. Here, we detail the implementation of an alternative recursive evaluation algorithm², reminiscent of dynamic programming concepts, which significantly reduces the number of arithmetic operations at the cost of introducing additional temporary storage requirements.

The idea is to express the basis functions of higher correlation order in terms of a product of just two basis functions of a lower correlation order. Consider the many-body basis in terms of \mathbf{v} tuples indexing into the atomic base A . We say that a tuple \mathbf{v} of length v has a decomposition $\mathbf{v} \equiv \mathbf{v}' \cup \mathbf{v}''$, where $\mathbf{v}', \mathbf{v}''$ have lengths v', v'' , if the tuples $(v_1, \dots, v_{v'})$ and $(v', \dots, v', v'', \dots, v'')$ agree up to permutations. In this case, we can write:

$$\mathbf{A}_{iV} = \mathbf{A}_{iV'} \mathbf{A}_{iV''}; \quad (29)$$

that is, the basis function \mathbf{A}_{iV} can be computed with a single product instead of $v - 1$ products, while its adjoint requires no additional products.

Nigam et al.'s NICE method⁵⁵ can be understood as an alternative recursive evaluation of ACE. It exploits the recurrence relations for coupling coefficients, which we also employ in the construction of the fully symmetric ACE basis \mathbf{B} in ref. ². In contrast our recursive evaluator uses recursion to optimize the evaluation of the \mathbf{A} basis that is only permutation-invariant, but significantly faster to construct than the \mathbf{B} basis—with the rotational invariance encoded in the basis coefficients.

A key subtlety must be addressed before putting this into practise: due to the constraints that $\sum_t m_t = 0$ and $\sum_t l_t$ even, not all basis functions have a decomposition (29) that respects those constraints. For example, if $\mathbf{m} = (1, 0, -1)$ then we may decompose it as $(1, -1) \cup (0)$, but if $\mathbf{m} = (2, -1, -1)$ then no such decomposition exists. To overcome this we add “artificial” basis functions to the model supplied with zero coefficients. A simple but seemingly effective heuristic how to achieve this efficiently is described in the following. The result of this construction is a directed acyclic graph:

$$\mathcal{G} = \{\mathbf{v} \equiv \mathbf{v}' \cup \mathbf{v}''\}, \quad (30)$$

where each node \mathbf{v} represents a basis function supplied with coefficient $\tilde{\mathbf{c}}_{iV}^{(p)}$ with exactly two incoming edges $(\mathbf{v}', \mathbf{v}), (\mathbf{v}'', \mathbf{v})$ and arbitrarily many (possibly zero) outgoing edges. The values of the coefficients are readily obtained from the canonical basis representation.

To construct \mathcal{G} we first insert the atomic base $\{A_{iV}\}$ represented by its indices $\{V\}$ into the graph as root nodes. Then, with increasing correlation order we insert the nodes \mathbf{v} using the following recursive algorithm.

Algorithm 6. Insert node \mathbf{v} into graph \mathcal{G}

```

1: for all decompositions  $\mathbf{v} \equiv \mathbf{v}' \cup \mathbf{v}''$  do
2:   if  $\mathbf{v}', \mathbf{v}'' \in \mathcal{G}$  then
3:      $\mathcal{G} \leftarrow \mathcal{G} \cup \{\mathbf{v} \equiv \mathbf{v}' \cup \mathbf{v}''\}$ 
4:   return
5: end if
6: end for
7: identify  $\mathbf{v} = (v_1, \dots) \cup \mathbf{v}'$  where  $v_1$  is maximal
8: recursively insert  $\mathbf{v}'$  into  $\mathcal{G}$  with zero coefficients
9: recursively insert  $\mathbf{v}$  into  $\mathcal{G}$ 

```

This simple heuristic already leads to excellent performance, as we report at the end of this section, but further optimizations may be possible to the graph aiming to minimize the number of artificial nodes inserted into the graph and limiting the additionally required memory access.

To evaluate the properties $\varphi_i^{(p)}$ we first apply Algorithm 1 to obtain the atomic base $A_{iV} \equiv A_{i\mu n l m}$. Next, we can traverse the graph taking care to only evaluate basis functions whose parents have already been evaluated (this is implicitly assumed), e.g., by looping with increasing correlation order.

Algorithm 2R. Recursive evaluation of properties

```

1:  $\varphi_i^{(p)} = 0$ 
2: for  $\mathbf{v} \equiv \mathbf{v}' \cup \mathbf{v}'' \leftarrow \mathcal{G}$  do
3:    $\mathbf{A}_{iV} = \mathbf{A}_{iV'} \mathbf{A}_{iV''}$ 
4:   for  $p$  do
5:      $\varphi_i^{(p)} += \tilde{\mathbf{c}}_{iV}^{(p)} \text{Re}(\mathbf{A}_{iV})$ 
6:   end for
7: end for

```

Only the values \mathbf{A}_{iV} corresponding to interior nodes \mathbf{v} (i.e., nodes that have at least one child) must be stored, while those corresponding to leaf nodes (i.e., nodes without any child) are only required locally to update $\varphi_i^{(p)}$.

To evaluate the adjoints $\omega_{iV} \equiv \omega_{i\mu n l m}$ we use the observation that:

$$\partial \mathbf{A}_{iV} = \partial (\mathbf{A}_{iV'} \mathbf{A}_{iV''}) = \mathbf{A}_{iV'} \partial \mathbf{A}_{iV''} + \mathbf{A}_{iV''} \partial \mathbf{A}_{iV'}, \quad (31)$$

where ∂ is a differential operator, and hence:

$$\frac{\partial \mathcal{F}(\varphi_i^{(1)}, \dots, \varphi_i^{(p)})}{\partial \mathbf{A}_{iV}} \partial \mathbf{A}_{iV} = \sum_p \frac{\partial \mathcal{F}}{\partial \varphi_i^{(p)}} \tilde{\mathbf{c}}_{iV}^{(p)} \partial \mathbf{A}_{iV} = (\theta_v \mathbf{A}_{iV'}) \partial \mathbf{A}_{iV''} + (\theta_v \mathbf{A}_{iV''}) \partial \mathbf{A}_{iV'}. \quad (32)$$

This shows that the adjoint of \mathbf{A}_{iV} can be propagated to the adjoints of the two parents. This immediately leads to the following reverse mode

differentiation algorithm, which computes adjoints $\omega_{\mathbf{v}}$ for all basis functions $\mathbf{A}_{\mathbf{v}}$ (or at least those corresponding to interior nodes of the graph). However, only the adjoints for root nodes, $\omega_{\mathbf{v}} = \omega_{\mu_{\text{in}} \mu_{\text{out}}}$, are eventually used to assemble the forces (Eq. (16)). The traversal of the graph must now be done in reverse order, that is, a node \mathbf{v} may only be visited once all of its children have been visited, for example, by traversing in reverse correlation order.

To conclude, we now use Algorithm 5 to evaluate the forces.

Algorithm 3R. Recursive evaluation of adjoints ω

```

1:  $\omega_{\mathbf{v}} = 0$ 
2: for  $\mathbf{v} \equiv \mathbf{v}' \cup \mathbf{v}'' \leftarrow \mathcal{G}$ , in reverse order do
3:    $\theta_{\mathbf{v}} = \sum_p \partial \mathcal{F} / \partial \phi_i^{(p)} \bar{c}_{\mu, \mathbf{v}}^{(p)}$ 
4:    $\omega_{\mathbf{v}} += \theta_{\mathbf{v}}$ 
5:    $\omega_{\mathbf{v}'} += \omega_{\mathbf{v}} \mathbf{A}_{\mathbf{v}'}$ 
6:    $\omega_{\mathbf{v}''} += \omega_{\mathbf{v}} \mathbf{A}_{\mathbf{v}''}$ 
7: end for

```

The forward pass, Algorithm 2R, requires $1 + P$ multiplications and $5 + P$ memory access operations at each iteration. The backward pass, Algorithm 3R, requires $2 + P$ multiplications and $7 + P$ memory access operations. In particular, the cost is (seemingly) independent of the correlation order of each basis functions. The overall cost is given by:

$$O(N \cdot \#A) + O(\#G \cdot P), \quad (33)$$

i.e., it scales linearly in the number of neighbors N and the number of nodes in the graph. The first part for setting up the atomic base is unaffected by the recursive evaluator.

Comparing the cost between the two approaches is difficult since we have no estimates on the number of artificial nodes that must be inserted into the graph. In practise we observe that there are always more leaf nodes than interior nodes, which means that relatively few artificial nodes are inserted and hence the recursive algorithm is significantly faster for large basis sets and high correlation order, but roughly comparable for small basis sets and low correlation order.

For the Cu potential with a smaller number of basis functions the timing decreases from 0.43 to 0.32 ms/atom/MD-step when the recursive evaluator is used. The Si potential, employing more basis functions, has a more pronounced speed-up from 1.84 to 0.80 ms/atom/MD-step.

DATA AVAILABILITY

The potentials files for Cu and Si and reference data for Cu are available at <https://doi.org/10.5281/zenodo.4734036>.

CODE AVAILABILITY

The PACE implementation described here is distributed with LAMMPS as the USER-PACE package. https://docs.lammps.org/Packages_details.html#pkg-user-pace. The USER-PACE package can also be accessed directly from <https://github.com/ICAMS/lammps-user-pace>.

Received: 31 January 2021; Accepted: 19 May 2021;

Published online: 28 June 2021

REFERENCES

1. Drautz, R. Atomic cluster expansion for accurate and transferable interatomic potentials. *Phys. Rev. B* **99**, 014104 (2019).
2. Dussan, G. et al. Atomic cluster expansion: Completeness, efficiency and stability. Preprint at <https://arxiv.org/abs/1911.03550> (2020).
3. 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).
4. Finnis, M. W. & Sinclair, J. E. A simple empirical N-body potential for transition metals. *Philos. Mag. A* **50**, 45–55 (1984).
5. Pettifor, D. G. & Oleinik, I. I. Bounded analytic bond-order potentials for σ and π bonds. *Phys. Rev. Lett.* **84**, 4124–4127 (2000).

6. Drautz, R. & Pettifor, D. G. Valence-dependent analytic bond-order potential for magnetic transition metals. *Phys. Rev. B* **84**, 214114 (2011).
7. Shapeev, A. V. Moment tensor potentials: a class of systematically improvable interatomic potentials. *Mult. Model. Simul.* **14**, 1153–1173 (2016).
8. Seko, A., Togo, A. & Tanaka, I. Group-theoretical high-order rotational invariants for structural representations: application to linearized machine learning interatomic potential. *Phys. Rev. B* **99**, 214108 (2019).
9. Thompson, A., Swiler, L., Trott, C., Foiles, S. & Tucker, G. Spectral neighbor analysis method for automated generation of quantum-accurate interatomic potentials. *J. Comp. Phys.* **285**, 316–330 (2015).
10. van der Oord, C., Dussan, G., Csányi, G. & Ortner, C. Regularised atomic body-ordered permutation-invariant polynomials for the construction of interatomic potentials. *Mach. Learn.: Sci. Technol.* **1**, 015004 (2020).
11. Behler, J. Atom-centered symmetry functions for constructing high-dimensional neural network potentials. *J. Chem. Phys.* **134**, 074106 (2011).
12. Bartók, A. P., Kondor, R. & Csányi, G. On representing chemical environments. *Phys. Rev. B* **87**, 184115 (2013).
13. Drautz, R. Atomic cluster expansion of scalar, vectorial, and tensorial properties including magnetism and charge transfer. *Phys. Rev. B* **102**, 024104 (2020).
14. Plimpton, S. Fast parallel algorithms for short-range molecular dynamics. *J. Comp. Phys.* **117**, 1–19 (1995).
15. Zuo, Y. et al. Performance and cost assessment of machine learning interatomic potentials. *J. Phys. Chem. A* **124**, 731–745 (2020).
16. 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).
17. Behler, J. & Parrinello, M. Generalized neural-network representation of high-dimensional potential-energy surfaces. *Phys. Rev. Lett.* **98**, 146401 (2007).
18. Drautz, R., Fähnle, M. & Sanchez, J. M. General relations between many-body potentials and cluster expansions in multicomponent systems. *J. Phys.: Condens. Matter* **16**, 3843–3852 (2004).
19. Drautz, R. et al. Analytic bond-order potential for predicting structural trends across the sp-valent elements. *Phys. Rev. B* **72**, 144105 (2005).
20. Stillinger, F. H. & Weber, T. A. Computer simulation of local order in condensed phases of silicon. *Phys. Rev. B* **31**, 5262–5271 (1985).
21. Tersoff, J. Empirical interatomic potential for silicon with improved elastic properties. *Phys. Rev. B* **38**, 9902–9905 (1988).
22. Bartók, A. P., Kermode, J., Bernstein, N. & Csányi, G. Machine learning a general-purpose interatomic potential for silicon. *Phys. Rev. X* **8**, 041048 (2018).
23. Mishin, Y., Mehl, M., Papaconstantopoulos, D., Voter, A. & Kress, J. Structural stability and lattice defects in copper: ab initio, tight-binding, and embedded-atom calculations. *Phys. Rev. B* **63**, 224106 (2001).
24. Li, X.-G. et al. Quantum-accurate spectral neighbor analysis potential models for Ni-Mo binary alloys and fcc metals. *Phys. Rev. B* **98**, 094104 (2018).
25. Perdew, J. P., Burke, K. & Ernzerhof, M. Generalized gradient approximation made simple. *Phys. Rev. Lett.* **77**, 3865–3868 (1996).
26. Blum, V. et al. Ab initio molecular simulations with numeric atom-centered orbitals. *Comput. Phys. Commun.* **180**, 2175–2196 (2009).
27. Havu, V., Blum, V., Havu, P. & Scheffler, M. Efficient O(N) integration for all-electron electronic structure calculation using numerically tabulated basis functions. *J. Comp. Phys.* **228**, 8367–8379 (2009).
28. Janssen, J. et al. pyiron: an integrated development environment for computational materials science. *Comput. Mater. Sci.* **163**, 24–36 (2019).
29. Clark, S. J. et al. First principles methods using CASTEP. *Z. Kristallogr. Cryst. Mater.* **220**, 567–570 (2005).
30. Apopstol, F. & Mishin, Y. Interatomic potential for the Al-Cu system. *Phys. Rev. B* **83**, 054116 (2011).
31. Etesami, S. & Asadi, E. Molecular dynamics for near melting temperatures simulations of metals using modified embedded-atom method. *J. Phys. Chem. Solids* **112**, 61–72 (2018).
32. Tersoff, J. New empirical approach for the structure and energy of covalent systems. *Phys. Rev. B* **37**, 6991–7000 (1988).
33. Du, Y., Lenosky, T., Hennig, R., Goedecker, S. & Wilkins, J. Energy landscape of silicon tetra-interstitials using an optimized classical potential. *Phys. Stat. Sol. B* **248**, 2050–2055 (2011).
34. Starikov, S., Gordeev, I., Lysogorskiy, Y., Kolotova, L. & Makarov, S. Optimized interatomic potential for study of structure and phase transitions in Si-Au and Si-Al systems. *Comput. Mater. Sci.* **184**, 109891 (2020).
35. 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).
36. Luo, W., Roundy, D., Cohen, M. & Morris, J. Ideal strength of bcc molybdenum and niobium. *Phys. Rev. B* **66**, 094110 (2002).

37. Zhu, L.-F., Grabowski, B. & Neugebauer, J. Efficient approach to compute melting properties fully from ab initio with application to Cu. *Phys. Rev. B* **96**, 224202 (2017).
38. Zheng, H. et al. Grain boundary properties of elemental metals. *Acta Mater.* **186**, 40–49 (2020).
39. Nguyen-Manh, D., Pettifor, D. G. & Vitek, V. Analytic environment-dependent tight-binding bond integrals: application to MoSi₂. *Phys. Rev. Lett.* **85**, 4136–4139 (2000).
40. Ma, P.-W. & Dudarev, S. L. Nonuniversal structure of point defects in face-centered cubic metals. *Phys. Rev. Mat.* **5**, 013601 (2021).
41. Connétable, D., Andrieu, E. & Monceau, D. First-principles nickel database: energetics of impurities and defects. *Comput. Mater. Sci.* **101**, 77–87 (2015).
42. Cogollo-Olivo, B. H., Seriani, N. & Montoya, J. A. Unbiased structural search of small copper clusters within DFT. *Chem. Phys.* **461**, 20–24 (2015).
43. Ono, S. Dynamical stability of two-dimensional metals in the periodic table. *Phys. Rev. B* **102**, 165424 (2020).
44. Stillinger, F. H. & Weber, T. A. Erratum: computer simulation of local order in condensed phases of silicon. *Phys. Rev. B* **33**, 1451–1451 (1986).
45. Justo, J. F., Bazant, M. Z., Kaxiras, E., Bulatov, V. V. & Yip, S. Interatomic potential for silicon defects and disordered phases. *Phys. Rev. B* **58**, 2539–2550 (1998).
46. Baskes, M. I. Modified embedded-atom potentials for cubic materials and impurities. *Phys. Rev. B* **46**, 2727–2742 (1992).
47. Porezag, D., Frauenheim, T., Köhler, T., Seifert, G. & Kaschner, R. Construction of tight-binding-like potentials on the basis of density-functional theory: application to carbon. *Phys. Rev. B* **51**, 12947–12957 (1995).
48. Buehler, M. J., van Duin, A. C. T. & Goddard, W. A. Multiparadigm modeling of dynamical crack propagation in silicon using a reactive force field. *Phys. Rev. Lett.* **96**, 095505 (2006).
49. Pickard, C. J. & Needs, R. J. Ab initio random structure searching. *J. Phys. Condens. Matter* **23**, 053201 (2011).
50. Okada, Y. & Tokumaru, Y. Precise determination of lattice parameter and thermal expansion coefficient of silicon between 300 and 1500 K. *J. Appl. Phys.* **56**, 314–320 (1984).
51. Laaziri, K. et al. High resolution radial distribution function of pure amorphous silicon. *Phys. Rev. Lett.* **82**, 3460–3463 (1999).
52. Goedecker, S., Deutsch, T. & Billard, L. A fourfold coordinated point defect in silicon. *Phys. Rev. Lett.* **88**, 235501 (2002).
53. Pickard, C. J. & Needs, R. J. High-pressure phases of silane. *Phys. Rev. Lett.* **97**, 045504 (2006).
54. Barzilai, J. & Borwein, J. M. Two-point step size gradient methods. *IMA J. Numer. Anal.* **8**, 141–148 (1988).
55. Nigam, J., Pozdnyakov, S. & Ceriotti, M. Recursive evaluation and iterative contraction of n-body equivariant features. *J. Chem. Phys.* **53**, 121101 (2020).
56. Ledbetter, H. Elastic constants of polycrystalline copper at low temperatures. relationship to single-crystal elastic constants. *phys. stat. sol. (a)* **66**, 477–484 (1981).
57. Siegel, R. Vacancy concentrations in metals. *J. Nucl. Mater.* **69**, 117–146 (1978).
58. Ehrhart, P. *Atomic Defects in Metals* (Landolt-Bornstein, New Series, 1991).
59. Ullmaier, H. *Properties and Interaction of Atomic Defects in Metals and Alloys Vol. 25*, 88 (Landolt-Bornstein, New Series, Group III, 1991).
60. Chekhovskoi, V. Y., Tarasov, V. D. & Gusev, Y. V. Calorific properties of liquid copper. *High Temp.* **38**, 394–399 (2000).
61. Wang, K. & Reeber, R. R. Thermal expansion of copper. *High Temp. Mat. Sci.* **35**, 181–186 (1996).

ACKNOWLEDGEMENTS

The authors acknowledge helpful discussions with Marc Cawkwell. R.D. acknowledges funding through the German Science Foundation (DFG), project number 405621217. Sandia National Laboratories is a multimission laboratory managed and operated by National Technology & Engineering Solutions of Sandia, LLC, a wholly owned subsidiary of Honeywell International Inc., for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-NA0003525.

AUTHOR CONTRIBUTIONS

Y.L., C.O. and R.D. led the implementation of PACE. Y.L. and R.D. carried out the DFT calculations for Cu. The parameterization and testing for Cu was largely done by Y.L., M.M., A.B., S.M. and R.D., for Si by C.v.d.O., C.O. and G.C. Y.L. and A.T. ensured compatibility with LAMMPS. R.D. wrote the initial version of the manuscript and figures were generated by Y.L. and C.v.d.O. All authors contributed to the implementation of PACE, the parameterization and testing of the ACE for Cu and Si and to editing and discussing the manuscript.

FUNDING

Open Access funding enabled and organized by Projekt DEAL.

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/s41524-021-00559-9>.

Correspondence and requests for materials should be addressed to R.D.

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