



CHORUS

This is the accepted manuscript made available via CHORUS. The article has been published as:

Efficient and accurate machine-learning interpolation of atomic energies in compositions with many species

Nongnuch Artrith, Alexander Urban, and Gerbrand Ceder

Phys. Rev. B **96**, 014112 — Published 21 July 2017

DOI: [10.1103/PhysRevB.96.014112](https://doi.org/10.1103/PhysRevB.96.014112)

Efficient and Accurate Machine-Learning Interpolation of Atomic Energies in Compositions with Many Species

Nongnuch Artrith,^{1,*} Alexander Urban,¹ and Gerbrand Ceder^{1,2,†}

¹*Department of Materials Science and Engineering, University of California, Berkeley, CA, USA*

²*Materials Science Division, Lawrence Berkeley National Laboratory, Berkeley, CA, USA*

(Dated: July 7, 2017)

Machine-learning potentials (MLPs) for atomistic simulations are a promising alternative to conventional classical potentials. Current approaches rely on descriptors of the local atomic environment with dimensions that increase quadratically with the number of chemical species. In this article, we demonstrate that such a scaling can be avoided in practice. We show that a mathematically simple and computationally efficient descriptor with constant complexity is sufficient to represent transition-metal oxide compositions and biomolecules containing 11 chemical species with a precision of around 3 meV/atom. This insight removes a perceived bound on the utility of MLPs and paves the way to investigate the physics of previously inaccessible materials with more than ten chemical species.

Atomic interaction potentials based on the interpolation of *first-principles* calculations with machine-learning algorithms have the potential to enable efficient linear-scaling atomistic simulations with an accuracy that is close to the reference method [1–4]. Such machine-learning potentials (MLPs) establish a relationship between a unique descriptor and the total or atomic energy using, e.g., artificial neural networks (ANNs) [5] or Gaussian process regression (GPR, Kriging) [6]. However, the combined space of atomic coordinates and chemical species grows rapidly with the number of chemical species, resulting in a formal corresponding growth of the descriptor complexity and thus the complexity of the MLP. This scaling has so far limited current MLP approaches to compositions with only a few chemical species [7–10] or atomic structures [11]. Overcoming this limitation is a very active field of research [12, 13].

In this article we demonstrate that the computational complexity of MLPs does not necessarily grow with the number of chemical species, so that MLPs for materials with ten or more chemical species are in principle feasible and computationally efficient. We show that, contrary to intuition and common belief, the same model complexity that is optimal for a ternary material is also sufficient to describe a system with 11 chemical species (**Fig. 1**). To illustrate these concepts, we consider two different materials classes of practical relevance: cation-disordered lithium transition-metal (TM) oxides, which have recently attracted interest as high-energy-density cathode materials for Li-ion batteries [14, 15], and proteinogenic amino acids, i.e., the building blocks of proteins and their complexes with divalent cations [16, 17]. We show that both of these high-dimensional materials systems can be accurately modeled using MLPs based on a mathematically simple and computationally efficient descriptor with constant complexity that we will introduce in the following.

In the present work, we focus on MLPs that express the total structural energy as the sum of atomic energy con-

tributions and are in this respect similar to other many-body potentials such as embedded atom models [18, 19]. However, unlike conventional potentials, the atomic energy is not confined to a rigid functional form, but is represented by a flexible non-linear machine-learning model that is trained to a descriptor of the *local atomic environment*. In this context, the local atomic environment $\sigma_i^{R_c} \subset \sigma$ of an atom i in a structure σ is defined as the *local structure* given by the set of coordinates $\{\mathbf{R}_1, \mathbf{R}_2, \dots\}$ of all atoms within a cutoff distance R_c from atom i and the *local composition*, i.e., the corresponding chemical species $\{t_1, t_2, \dots\}$. To be physically meaningful and transferable between equivalent structures, the descriptor needs to be invariant with respect to translation and rotation of the structure and the exchange of equivalent atoms. Several transformations for $\sigma_i^{R_c}$ into invariant representations $\tilde{\sigma}_i^{R_c}$ have been proposed in the literature [20–26], and the most commonly used methods for MLPs are the *symmetry functions* by Behler and Parrinello (BP) [2, 20] and the *smooth overlap of atomic positions* (SOAP) approach by Bartók, Kondor, and Csányi [13, 21, 27]. With an invariant descriptor $\tilde{\sigma}_i^{R_c}$, the total MLP energy of a structure σ can then be expressed as

$$E(\sigma) = \sum_i^{\text{atoms}} \text{MLP}_{t_i}(\tilde{\sigma}_i^{R_c}) \quad .$$

Our approach draws inspiration from the strength of the established descriptor methods but explicitly maintains the distinction between *local structure* and *composition* by using two sets of invariant coordinates, $\{\mathbf{R}\}\tilde{\sigma}_i^{R_c}$ and $\{t\}\tilde{\sigma}_i^{R_c}$, that separately encode the atomic positions and species. The union of both sets, $\tilde{\sigma}_i^{R_c} = \{\mathbf{R}\}\tilde{\sigma}_i^{R_c} \cup \{t\}\tilde{\sigma}_i^{R_c}$, is used as a combined descriptor for an ANN-based MLP (ANN potential). As *structural* descriptor $\{\mathbf{R}\}\tilde{\sigma}_i^{R_c}$ we choose the expansion coefficients of the radial (bond length) and angular (bond angle) distribution

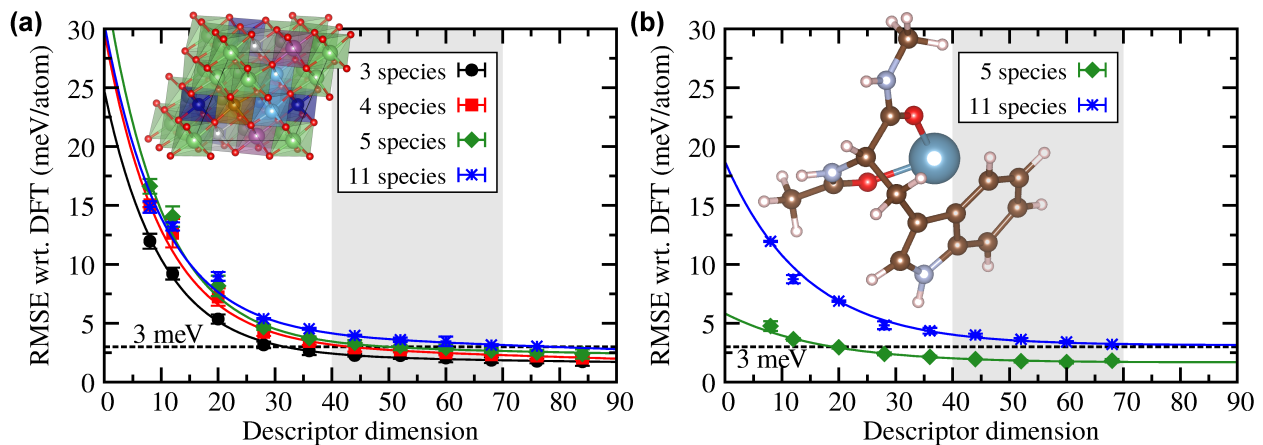


FIG. 1. Precision of artificial neural network (ANN) potentials as function of the dimension of the descriptor used to represent the local atomic environment. **(a)** Root mean squared error (RMSE) of the ANN potential energies relative to their DFT references for LiMO_2 systems with increasing number of chemical species: 3 species ($M = \text{Ti}$; black circles), 4 species ($M = \text{Ti}, \text{Ni}$; red squares), 5 species ($M = \text{Ti}, \text{Mn}, \text{Ni}$; green diamonds), and 11 species ($M = \text{Sc}, \text{Ti}, \text{V}, \text{Cr}, \text{Mn}, \text{Fe}, \text{Co}, \text{Ni}, \text{Cu}$; blue stars). The unit cell of a representative LiMO_2 structure from the data set is shown as inset. **(b)** An equivalent analysis for a data set with conformations of the 20 proteinogenic amino acids (5 chemical species: H, C, N, O, S; green diamonds) and their complexes with the divalent cations Ba^{2+} , Ca^{2+} , Cd^{2+} , Hg^{2+} , Pb^{2+} , Sr^{2+} (in total 11 species; blue stars). The inset shows one conformation of a tryptophan dipeptide complex with Ca^{2+} . Generally, the RMSE was evaluated after 3,000 training iterations, except for the two 11-species systems for which 5,000 iterations were required. The error bars indicate the standard deviation of three independently trained ANN potentials, the gray region highlights descriptors that result in essentially converged ANN potentials with RMSE values around 3 meV/atom, and the lines are meant to guide the eye.

functions in a complete basis set $\{\phi_\alpha\}$,

$$\text{RDF}_i(r) = \sum_{\alpha} c_{\alpha}^{(2)} \phi_{\alpha}(r) \quad \text{for } 0 \leq r \leq R_c \quad (1)$$

$$\text{ADF}_i(\theta) = \sum_{\alpha} c_{\alpha}^{(3)} \phi_{\alpha}(\theta) \quad \text{for } 0 \leq \theta \leq \pi \quad , \quad (2)$$

and the *compositional* descriptor $\{^t\}_{\sigma_i}^{\sim R_c}$ is given by the expansion coefficients of the same distribution functions but with atomic contributions that are weighted differently for each chemical species. The RDF and ADF obey the invariants of the atomic energy, and basing the descriptor on an expansion in a complete basis set allows its systematic refinement by converging the number of basis functions. We implemented the descriptor into the free and open-source *atomic energy network* package [28].

In general, multi-layer ANNs can reproduce any function with arbitrary precision [29]. However, the resolution of the invariant descriptor determines the maximal precision with which an ANN potential can resolve the chemical space of a given material. To determine the resolution of our combined descriptor, we trained ANN potentials to extensive reference data sets with different numbers of chemical species. We consider the resolution satisfactory if the ANN potential can reproduce the reference energies of our data sets with a precision of ~ 3 meV/atom, which is the order of magnitude of the noise in our reference data.

Figure 1a shows the precision that can be achieved in representing Li-TM oxides with different numbers of

TM species using ANN potentials based on the combined descriptor with different numbers of basis functions. The reference set for the ANN potential training comprised Hubbard-U corrected [30–32] density-functional theory (DFT) energies and optimized structures of 16,047 LiMO_2 configurations in the rocksalt structure with different compositions based on nine TMs (Sc, Ti, V, Cr, Mn, Fe, Co, Ni, and Cu) and cation arrangements with up to 36 atoms. For all DFT+U calculations we employed the PBE exchange-correlation functional [33] with projector-augmented wave [34] pseudopotentials as implemented in VASP [35, 36]. DFT energies and atomic forces were converged to 0.05 meV per atom and 50 meV/Å, respectively, gamma-centered k-point meshes with a density of 1000 divided by the number of atoms were used, and the plane-wave cutoff was 520 eV. VASP input files were generated using the pymatgen software with default parameters [37]. Structures with up to 5 chemical species were generated by systematic enumeration, and random atomic configurations were generated for compositions with 6–11 chemical species. Further information about the generation of these reference structures, the parameters of our DFT calculations, and the architecture of the ANNs are given in the electronic Supplemental Material (SM) [38].

As seen in **Fig. 1a**, the ANN potentials achieve a root mean squared error (RMSE) of ~ 3 meV/atom relative to the DFT reference energies with a descriptor dimension of 44 (i.e., 22 basis functions). Note that, for

the present work, we employed the same number of basis functions for the radial and angular expansion (i.e., 11 each), though this is not a general requirement of the methodology. Increasing the descriptor dimension beyond 52 or 60 results in a minor additional reduction of the RMSE at the cost of significantly increased computational effort. We emphasize that this RMSE is purely a quality measure of the descriptor precision and does not reflect the accuracy of the ANN potentials in simulations, which would have to be carefully validated separately.

The RMSE was evaluated after 3,000 training iterations using the LM-BFGS method [39, 40], however, with increasing number of species and increasing descriptor size the required number of training iterations to achieve convergence generally also increases. Thus, the ANN potentials for 11 chemical species and descriptor dimensions above 40 have not converged after 3,000 iterations, and the RMSEs after 5,000 iterations are shown in **Fig. 1**. The unconverged RMSE after 3,000 training iterations is shown in **Fig. S2** in the SM [38].

Remarkably, the optimal descriptor dimension is essentially independent of the number of chemical species in the composition, and a descriptor dimension of 44 is sufficient to capture the structural and chemical features of the distinct atomic configurations in the LiMO_2 data set with up to 11 chemical species.

Figure 1b shows the equivalent analysis for the first-principles energies and structures of 45,892 conformations of the proteinogenic amino acids (5 chemical species: H, C, N, O, and S) and their complexes with the six divalent cations Ba^{2+} , Ca^{2+} , Cd^{2+} , Hg^{2+} , Pb^{2+} , Sr^{2+} (a total of 11 chemical species) by Ropo, Schneider, Baldauf, and Blum [16] based on DFT calculations (PBE+TS-vdW [41]) using the FHI-AIMS package [42]. This data set was compiled specifically for the parametrization of atomic potentials and thoroughly samples the relevant conformational space [16], an important first step towards improved force fields for proteins [43]. The high precision of the ANN potentials with an RMSE of ~ 3 meV/atom for 5 and 11 chemical species indicates that our combined descriptor is not limited to crystal structures with similar atomic positions, but is also suitable to distinguish between continuous atomic arrangements.

To understand the significance of these observations, we first describe the details of the structural and compositional descriptor. We begin by expressing the atom-centered radial and angular distribution functions of Eqs. (1) and (2) in terms of discrete delta functions centered at the bond lengths between atoms j and the central atom i , $R_{ij} = \|\mathbf{R}_j - \mathbf{R}_i\|$, and the bond angle




FIG. 2. (a) Discrete atom-centered radial distribution function (RDF_i) for a lithium site in a structure with composition $\text{Li}_2\text{MnNiO}_4$ (black lines) and the cosine cutoff function f_c for a cutoff radius of $R_c = 8 \text{ \AA}$. (b) Convolution of the RDF of (a) with a Gaussian function with a width of 0.2 \AA (black line) and the reconstructed RDF from a Chebyshev expansion with a radial order $N_2 = 50$ (orange line). (c) Same as (b), but with a Gaussian width of 0.1 \AA and an expansion order of $N_2 = 150$.

$$\theta_{ijk} = \angle(\mathbf{R}_j - \mathbf{R}_i, \mathbf{R}_k - \mathbf{R}_i)$$

$$\text{RDF}_i(r) = \sum_{\mathbf{R}_j \in \sigma_i^{R_c}} \delta(r - R_{ij}) f_c(R_{ij}) w_{t_j} \quad (3)$$

$$\text{ADF}_i(\theta) = \sum_{\mathbf{R}_j, \mathbf{R}_k \in \sigma_i^{R_c}} \delta(\theta - \theta_{ijk}) f_c(R_{ij}) f_c(R_{ik}) w_{t_j} w_{t_k} \quad (4)$$

where f_c is a cutoff function that smoothly goes to zero at R_c (in practice, we use $f_c(r) = 0.5[\cos(r \cdot \pi/R_c) - 1]$). The weights w_{t_j} and w_{t_k} are 1 for the *structural descriptor* $\{\mathbf{R}\} \tilde{\sigma}_i^{R_c}$ and take on species-dependent values for the *compositional descriptor* $\{t\} \tilde{\sigma}_i^{R_c}$. Here, we followed the (Ising-model) pseudo-spin convention commonly used for lattice models [44], i.e., $w_l = 0, \pm 1, \pm 2, \dots$ where 0 is omitted for even numbers of species. For the expansions Eqs. (1) and (2) we choose a complete orthonormal basis $\{\phi_\alpha\}$, i.e., $\int \phi_\alpha \bar{\phi}_\beta = 1$ if $\alpha = \beta$ and 0 else. With this choice, the expansion coefficients are given by

$$c_\alpha^{(2)} = \sum_{\mathbf{R}_j \in \sigma_i^{R_c}} \phi_\alpha(R_{ij}) f_c(R_{ij}) w_{t_j} \quad \text{and} \quad (5)$$

$$c_\alpha^{(3)} = \sum_{\mathbf{R}_j, \mathbf{R}_k \in \sigma_i^{R_c}} \phi_\alpha(\theta_{ijk}) f_c(R_{ij}) f_c(R_{ik}) w_{t_j} w_{t_k} \quad (6)$$

A derivation of Eqs. (5) and (6) can be found in the SM [38]. The expansions are truncated at finite radial and angular orders N_2 and N_3 that determine the dimension (i.e., the complexity) and

the resolution of the descriptor, i.e., $\{\mathbf{R}\}_{\tilde{\sigma}_i^{Rc}} = \{\{\mathbf{R}\}_{c_1^{(2)}}, \dots, \{\mathbf{R}\}_{c_{N_2}^{(2)}}, \{\mathbf{R}\}_{c_1^{(3)}}, \dots, \{\mathbf{R}\}_{c_{N_3}^{(3)}}\}$.

For this article, we employed the Chebyshev polynomials of the first kind as basis functions (see SM [38]), as they can be defined in terms of a recurrence relation that allows for highly efficient numerical evaluation of the function values and their derivatives. With this choice of basis functions, **Fig. 2** shows the RDF as reconstructed based on the structural expansion coefficients $\{\{\mathbf{R}\}_{c_\alpha^{(2)}}\}$ for two different orders ($N_2 = 50$ and $N_2 = 150$). From comparison with Gaussian convolutions of the discrete RDF, the radial resolution of the expansion order $N_2 = 150$ is around 0.1 Å. Atomic features on smaller scales may affect the shape of the RDF but do not give rise to distinct peaks. The expansion of the ADF is completely analogous.

We note that the radial and angular BP symmetry functions [2, 20] can be cast into the form of Eqs. (5) and (6) but are neither orthogonal nor systematically refinable. The relationship of our structural descriptor to the coefficients of a basis set expansion is, in turn, closer in spirit to the SOAP method [3, 21] which is based on the power spectrum of the atomic density of the local atomic environment. SOAP allows for a rigorous and systematic description of the local structure, which comes at the cost of an arithmetically (and computationally) more complex formalism. However, by limiting the descriptor to radial and angular contributions our method maintains the simple analytic nature of the BP approach that allows for a highly efficient numerical implementation and straightforward differentiation (which is required for the calculation of analytic forces and higher derivatives). Basing the radial and angular descriptors on an expansion in a complete basis set allows their systematic refinement in the spirit of the SOAP approach, though our approach is limited to two- and three-body interactions.

Also note that decomposing the local atomic environment into n -body contributions as done in our structural descriptor is an established and well-tested approach for lattice models such as the cluster expansion (CE) method [45, 46]. In CE models, the total configurational energy is expanded in a basis set consisting of site clusters ($s_i s_j \dots$) with increasing numbers of lattice sites s_α , i.e., point clusters, pairs, trimers, \dots , n -tuples. The site clusters form a complete basis set, and the configurational averages of all equivalent clusters (the cluster correlations) are the descriptor of the CE model. Unlike MLPs, the CE energy is a *linear* function of the descriptor. For the case of the continuous structural energy, Thompson *et al.* demonstrated that a linear potential based on SOAP (which also is a complete basis of the local structure) can achieve reasonable accuracy in practice if a sufficient number of basis functions is used [47].

However, the strength of non-linear machine-learning models is that they do not require mathematically com-




FIG. 3. Convergence of the root mean squared error (RMSE) of the predicted LiMO_2 (5 species, $M = \text{Ti, Mn, Ni}$) energy for a linear model (empty diamonds) and a non-linear ANN (filled diamonds) with the dimension of the combined descriptor. The gray region highlights descriptors that result in essentially converged ANN potentials with RMSE values around 3 meV/atom. The lines are meant to guide the eye.

plete descriptors as long as the descriptor is able to differentiate between all *relevant* samples. This property is exploited, for example, in the area of image recognition and text classification [48]. In practice this means that even an incomplete descriptor of the local atomic environment may be sufficient to construct a non-linear MLP if that descriptor is able to differentiate between all *relevant* local atomic structures, i.e., the descriptor does not have to resolve all hypothetically possible sets of three dimensional coordinates.

This behavior is exemplified in **Fig. 3**, which compares the precision of an ANN potential for the LiMO_2 data set with 5 chemical species (10,175 atomic configurations) with that of a linear energy model as function of the descriptor dimension. As seen in the figure, the ANN achieves an RMSE of ≤ 3 meV/atom with descriptor dimensions of 44 (22 basis functions) and larger. Comparison with **Fig. 2** shows that such a small basis set corresponds to a coarse representation of the RDF and the ADF, however, obviously this level of approximation is sufficient for the ANN potential to differentiate between all structural and compositional features in the reference set. This is not the case for the linear model whose RMSE is > 15 meV/atom even for a descriptor dimension of 84.

In conclusion, we showed that machine-learning potentials do not require (mathematically) complete descriptors of the local atomic environment to reproduce potential energy surfaces with high precision. With this insight, we devised a combined descriptor of the local atomic *structure* and *composition* whose complexity does not scale with the number of chemical species. The method is conceptually simple and allows for highly effi-

cient numerical implementations. The utility of the approach was demonstrated for two exemplary materials classes, lithium transition-metal oxides and amino acid complexes, each separately comprising compositions with 11 different chemical species. We showed that the potential energy landscape of both example systems can be represented with high precision by artificial neural network potentials using the combined descriptor achieving a resolution of around 3 meV/atom. Hence, machine-learning potentials are in practice not limited to compositions with small numbers of chemical species as previously argued in the literature and may be effective for the modeling of high-dimensional materials such as oxide solid solutions and peptide chains.

ACKNOWLEDGMENTS

This work was supported by the Office of Naval Research (ONR) under ONR award N00014-14-1-0444. This work used mainly the computational facilities of the Extreme Science and Engineering Discovery Environment (XSEDE), which is supported by National Science Foundation grant no. ACI-1053575. Additional computational resources from the University of California Berkeley, HPC Cluster (SAVIO) are also gratefully acknowledged.

* nartrith@berkeley.edu

† gceder@berkeley.edu

- [1] S. Lorenz, A. Groß, and M. Scheffler, *Chem. Phys. Lett.* **395**, 210 (2004).
- [2] J. Behler and M. Parrinello, *Phys. Rev. Lett.* **98**, 146401 (2007).
- [3] A. P. Bartók, M. C. Payne, R. Kondor, and G. Csányi, *Phys. Rev. Lett.* **104**, 136403 (2010).
- [4] M. Rupp, A. Tkatchenko, K.-R. Müller, and O. A. von Lilienfeld, *Phys. Rev. Lett.* **108**, 058301 (2012).
- [5] G. Montavon, G. B. Orr, and K.-R. Müller, eds., *Neural Networks: Tricks of the Trade (Second Edition)*, Lecture Notes in Computer Science, Vol. 7700 (Springer Berlin Heidelberg, 2012).
- [6] C. E. Rasmussen and C. K. I. Williams, *Gaussian Processes for Machine Learning* (MIT University Press Group Ltd, 2006).
- [7] N. Artrith, T. Morawietz, and J. Behler, *Phys. Rev. B* **83**, 153101 (2011).
- [8] N. Artrith and A. M. Kolpak, *Nano Lett.* **14**, 2670 (2014).
- [9] N. Artrith and A. M. Kolpak, *Comput. Mater. Sci.* **110**, 2028 (2015).
- [10] T. Morawietz, A. Singraber, C. Dellago, and J. Behler, *Proc. Natl. Acad. Sci. USA* **113**, 8368 (2016).
- [11] F. A. Faber, A. Lindmaa, O. A. von Lilienfeld, and R. Armiento, *Phys. Rev. Lett.* **117**, 135502 (2016).
- [12] H. Huo and M. Rupp, *arXiv* (2017), 1704.06439.
- [13] S. De, A. P. Bartók, G. Csányi, and M. Ceriotti, *Phys. Chem. Chem. Phys.* **18**, 13754 (2016).
- [14] J. Lee, A. Urban, X. Li, D. Su, G. Hautier, and G. Ceder, *Science* **343**, 519 (2014).
- [15] N. Yabuuchi, M. Takeuchi, M. Nakayama, H. Shiiba, M. Ogawa, K. Nakayama, T. Ohta, D. Endo, T. Ozaki, T. Inamasu, K. Sato, and S. Komaba, *Proc. Natl. Acad. Sci. USA* **112**, 76507655 (2015).
- [16] M. Ropo, M. Schneider, C. Baldauf, and V. Blum, *Sci. Data* **3**, 160009 (2016).
- [17] M. Ropo, V. Blum, and C. Baldauf, *Sci. Rep.* **6**, 35772 (2016).
- [18] M. S. Daw and M. I. Baskes, *Phys. Rev. B* **29**, 6443 (1984).
- [19] M. S. Daw, S. M. Foiles, and M. I. Baskes, *Mater. Sci. Rep.* **9**, 251 (1993).
- [20] J. Behler, *J. Chem. Phys.* **134**, 074106 (2011).
- [21] A. P. Bartók, R. Kondor, and G. Csányi, *Phys. Rev. B* **87**, 184115 (2013).
- [22] A. Sadeghi, S. A. Ghasemi, B. Schaefer, S. Mohr, M. A. Lill, and S. Goedecker, *J. Chem. Phys.* **139**, 184118 (2013).
- [23] K. T. Schütt, H. Glawe, F. Brockherde, A. Sanna, K. R. Müller, and E. K. U. Gross, *Phys. Rev. B* **89**, 205118 (2014).
- [24] J. Behler, *Int. J. Quantum Chem.* **115**, 1032 (2015).
- [25] F. Faber, A. Lindmaa, O. A. von Lilienfeld, and R. Armiento, *Int. J. Quantum Chem.* **115**, 1094 (2015).
- [26] O. A. von Lilienfeld, R. Ramakrishnan, M. Rupp, and A. Knoll, *Int. J. Quantum Chem.* **115**, 1084 (2015).
- [27] A. P. Bartók and G. Csányi, *Int. J. Quantum Chem.* **115**, 1051 (2015).
- [28] N. Artrith and A. Urban, *Comput. Mater. Sci.* **114**, 135 (2016).
- [29] K. Hornik, *Neural Networks* **4**, 251 (1991).
- [30] A. I. Liechtenstein, V. I. Anisimov, and J. Zaanen, *Phys. Rev. B* **52**, R5467 (1995).
- [31] V. I. Anisimov, F. Aryasetiawan, and A. I. Liechtenstein, *J. Phys.: Condens. Matter* **9**, 767 (1997).
- [32] A. Jain, G. Hautier, C. J. Moore, S. P. Ong, C. C. Fischer, T. Mueller, K. A. Persson, and G. Ceder, *Comput. Mater. Sci.* **50**, 2295 (2011).
- [33] J. Perdew, K. Burke, and M. Ernzerhof, *Phys. Rev. Lett.* **77**, 3865 (1996).
- [34] P. E. Blöchl, *Phys. Rev. B* **50**, 17953 (1994).
- [35] G. Kresse and J. Furthmüller, *Phys. Rev. B* **54**, 11169 (1996).
- [36] G. Kresse and J. Furthmüller, *Comput. Mater. Sci.* **6**, 15 (1996).
- [37] S. P. Ong, W. D. Richards, A. Jain, G. Hautier, M. Kocher, S. Cholia, D. Gunter, V. L. Chevrier, K. A. Persson, and G. Ceder, *Comput. Mater. Sci.* **68**, 314 (2013).
- [38] See Supplemental Material [url] for details of the reference data set, the architecture of the ANNs, a discussion of the scaling behavior of other local descriptors, and a derivation of the Chebyshev expansion coefficients. The Supplemental Material includes Refs. [49–53].
- [39] R. Byrd, P. Lu, J. Nocedal, and C. Zhu, *SIAM J. Sci. Comput.* **16**, 1190 (1995).
- [40] C. Zhu, R. H. Byrd, P. Lu, and J. Nocedal, *ACM T. Math Software* **23**, 550560 (1997).
- [41] A. Tkatchenko and M. Scheffler, *Phys. Rev. Lett.* **102**, 073005 (2009).
- [42] V. Blum, R. Gehrke, F. Hanke, P. Havu, V. Havu, X. Ren, K. Reuter, and M. Scheffler, *Comput. Phys.*

- Commun. **180**, 2175 (2009).
- [43] S. Piana, J. L. Klepeis, and D. E. Shaw, *Curr. Opin. Struct. Biol.* **24**, 98 (2014).
- [44] J. Sanchez, F. Ducastelle, and D. Gratias, *Physica A* **128**, 334 (1984).
- [45] D. D. Fontaine, *Solid State Phys.* **47**, 33176 (1994).
- [46] G. Ceder, *Comput. Mater. Sci.* **1**, 144 (1993).
- [47] A. Thompson, L. Swiler, C. Trott, S. Foiles, and G. Tucker, *J. Comput. Phys.* **285**, 316 (2014).
- [48] M. Rogati and Y. Yang, in *Proceedings of the eleventh international conference on Information and knowledge management - CIKM '02* (ACM Press, 2002).
- [49] A. Urban, I. Matts, A. Abdellahi, and G. Ceder, *Adv. Energy Mater.* **6**, 1600488 (2016).
- [50] N. Artrith, B. Hiller, and J. Behler, *physica status solidi (b)* **250**, 1191 (2013).
- [51] G. L. W. Hart and R. W. Forcade, *Phys. Rev. B* **77**, 224115 (2008).
- [52] G. L. W. Hart and R. W. Forcade, *Phys. Rev. B* **80**, 014120 (2009).
- [53] G. L. Hart, L. J. Nelson, and R. W. Forcade, *Comput. Mater. Sci.* **59**, 101 (2012).