

1 De novo exploration and self-guided learning of potential-energy 2 surfaces

3 Noam Bernstein,¹ Gábor Csányi,² and Volker L. Deringer^{2,3,*}

4 ¹*Center for Materials Physics and Technology,*
5 *U.S. Naval Research Laboratory, Washington, DC 20375, United States*

6 ²*Department of Engineering, University of Cambridge,*
7 *Cambridge CB2 1PZ, United Kingdom*

8 ³*Department of Chemistry, University of Cambridge,*
9 *Cambridge CB2 1EW, United Kingdom*

10 (Dated: September 6, 2019)

Interatomic potential models based on machine learning (ML) are rapidly developing as tools for materials simulations. However, because of their flexibility, they require large fitting databases that are normally created with substantial manual selection and tuning of reference configurations. Here, we show that ML potentials can be built in a largely automated fashion, exploring and fitting potential-energy surfaces from the beginning (*de novo*) within one and the same protocol. The key enabling step is the use of a configuration-averaged kernel metric that allows one to select the few most relevant and diverse structures at each step. The resulting potentials are accurate and robust for the wide range of configurations that occur during structure searching, despite only requiring a relatively small number of single-point DFT calculations on small unit cells. We apply the method to materials with diverse chemical nature and coordination environments, marking an important step toward the more routine application of ML potentials in physics, chemistry, and materials science.

INTRODUCTION

Atomic-scale modeling has become a cornerstone of scientific research. Quantum-mechanical methods, most prominently based on density-functional theory (DFT), describe the atomistic structures and physical properties of materials with high confidence¹; increasingly, they also make it possible to discover previously unknown crystal structures and synthesis targets². Still, quantum-mechanical materials simulations are severely limited by their high computational cost.

Machine learning (ML) has emerged as a promising approach to tackle this long-standing problem^{3–12}. ML-based interatomic potentials approximate the high-dimensional potential-energy surface (PES) by fitting to a reference database, which is usually computed at the DFT level. Once generated, ML potentials enable accurate simulations that are orders of magnitude faster than the reference method. They can solve challenging structural problems, as has been demonstrated for the atomic-scale deposition and growth of amorphous carbon films¹³, for proton-transfer mechanisms¹⁴ or dislocations in materials^{15,16}, involving thousands of atoms in the simulation. More recently, it was shown that ML potentials can be suitable tools for global structure searches targeting crystalline phases^{17–20}, clusters^{21–24}, and nanostructures²⁵.

Assembling the reference databases to which ML potentials are fitted is currently mostly a manual and laborious process, guided by the physical problem under study. The first artificial neural network (NN) type potential for materials³ was made by enumerating known crystal structures for silicon and used to describe high-pressure phase transitions^{26,27}. To incorporate vacancies, surfaces and so on, hierarchical databases for transition metals have been built that start with simple unit cells and gradually add relevant defect structures^{28,29}; liquid and amorphous materials can be described by iteratively grown databases that contain relatively small-sized MD snapshots^{30–33}. A “general-purpose” Gaussian Approximation Potential (GAP) ML model for elemental silicon was recently developed³⁴ which can describe crystalline phases with meV-per-atom accuracy, treat defects, cracks, and surfaces³⁵, and generate amorphous silicon structures in excellent agreement with experiment³⁶. Despite their success in achieving their stated goals, none of these potentials are expected to be even reasonable for crystal structures not included in their databases, say, hitherto undiscovered phases that only become stable at very high pressures.

42 In contrast, structure searching (that is, a global exploration of the PES) can be a suit-
 43 able approach for finding structures to be included in the training databases *in the first*
 44 *place*^{18–20,37}. The principal idea to explore configuration space with preliminary ML poten-
 45 tials is well established: since the first high-dimensional ML potentials have been made, it
 46 was shown how they can be refined by exploring unknown structures^{3,26,31}, and “on the fly”
 47 schemes were proposed to add required data while an MD simulation is being run^{5,38–40}. We
 48 have previously shown that the PES of boron can be iteratively sampled without prior knowl-
 49 edge of any crystal structures involved; we called the method “GAP-driven random structure
 50 searching” (GAP-RSS)¹⁸, reminiscent of the successful Ab Initio Random Structure Search-
 51 ing (AIRSS) approach^{41,42}. Subsequently we demonstrated, by way of an example, that the
 52 crystal structure of black phosphorus can be discovered by GAP-RSS within a few iterations,
 53 and we identified several previously unknown hypothetical allotropes of phosphorus¹⁹.

54 In the context of ML potential fitting, so-called “active learning” schemes which detect
 55 extrapolation (indicating when the potential moves away from known configurations) are
 56 currently receiving much attention. A query-by-committee active-learning approach was
 57 suggested in 2012 by Artrith and Behler: two NN potential fits are made to the same
 58 database, and if their prediction differs for a given (new) structure, this structure needs
 59 to be added to the database.⁴³ More recently, Jinnouchi et al. demonstrated how ab initio
 60 molecular dynamics (AIMD) simulations of specific systems can be sped up by active learning
 61 of the computed forces (in a modified GAP framework), using the predicted error of the
 62 Gaussian process to select new datapoints and to improve the speed of AIMD^{38,40}. In
 63 the context of structure exploration, Shapeev and co-workers employed Moment Tensor
 64 Potentials⁴⁴ with active learning⁴⁵ to generate ML potentials^{20,46}, and E and co-workers
 65 described a generalized active-learning scheme for deep neural network potentials⁴⁷. So far,
 66 these studies mainly focused on specific intermetallic systems, namely, Al–Mg⁴⁷ and Cu–Pd,
 67 Co–Nb–V, and Al–Ni–Ti⁴⁶, respectively. Furthermore, Podryabinkin et al. showed that their
 68 approach can identify existing and hypothetical boron allotropes²⁰.

69 In this work, we present an efficient and unified approach for generating reference
 70 databases for fitting ML potentials, exploring structural space from the beginning (*de*
 71 *novo*) by ML-driven searching and similarity measures, all without any prior knowledge of
 72 what structures are or are not relevant. In contrast with continuous active learning, our
 73 aim is to converge to a potential that can describe a wide range of configurations without

the need for additional fitting. We demonstrate the ability to cover a broad range of structures and chemistries, from graphite sheets to a densely packed transition metal. Our work provides conceptual insight into how computers can discover structural chemistry based on data and similarity measures alone, and it paves the way for a more routine application of ML potentials in materials discovery.

RESULTS

A unified framework for exploring and fitting structural space

The overarching aim is to construct a ML potential with minimal effort, both in terms of computational resources and in terms of input required from the user. In regard to the former, we use only single-point DFT computations to generate the fitting database¹⁸. In regard to the latter, we define general heuristics wherever possible, such that neither the protocol nor its parameters need to be manually tuned for a specific system. The ML architecture to which we couple our method is based on a hierarchical combination of two-, three-, and many-body descriptors³², and it uses GAP as the regressor⁴. The remaining two parameters that need to be set by the user are a “characteristic” distance and whether the material is primarily covalent or metallic. For the distance, we choose tabulated covalent (for C, B, and Si)⁴⁸ or metallic (for Ti) radii, depending on the nature of the system. These define the volume of the initial structures and the cutoffs for the ML descriptors (Methods section).

Our approach is based on an iterative cycle, as shown in the diagram in Fig. 1a. We generate ensembles of randomized structures as in the AIRSS framework^{41,42}, a structure-searching approach that is widely used in physics, chemistry, and materials science^{49–51}. In the first iteration, we generate 10,000 initial structures, from which we select the N most diverse ones using the leverage-score CUR algorithm⁵². In the context of PES models, the CUR algorithm was proposed⁵³ and then used^{29,32,34} for selection of sparse (representative) points for Gaussian process regression, and also proposed for selection of training configurations⁵⁴. The distance between candidate structures is quantified by the Smooth Overlap of Atomic Positions (SOAP) descriptor⁵⁵, which has been widely used in GAP fitting^{32,34} and in structural analysis^{56–58}. While SOAP is normally used to discriminate between pairs of

environments of individual atoms, we here use a *configuration-averaged* SOAP descriptor⁵⁷ that compares entire unit cells to one another (Methods section). We find that selecting the most representative structures is critical, because we can only evaluate a small number ($\ll 10,000$) with DFT. In addition, the starting configurations include dimers in vacuum at a wide range of bond lengths; this serves to capture the exchange repulsion at very short interatomic distances, and thereby to make the potentials more robust³².

With the starting configurations in hand, we perform single-point DFT computations and fit an initial potential to the resulting data; in subsequent iterations, we extend the database and thereby refine the potential¹⁸. In each iteration, we start from the same number of new random initial structures, and minimize their enthalpy using the GAP from the previous iteration. We then select the N most relevant and diverse configurations from the full set of configurations seen throughout the minimization trajectories, for which we employ a combination of Boltzmann-probability biased flat histogram sampling (to focus on low-energy structures) and leverage-score CUR (to select the most diverse structures among those), as illustrated in Fig. 1b. These selected configurations are evaluated using single-point DFT calculations and added to the fitting database.

The iterative procedure runs until the results are satisfactory. Here we terminate our searches after 2,500 DFT data points have been collected, and our results show this to be sufficient to discover and describe all structures discussed in the present work. Other quality criteria, such as based on the distribution of energies in the database¹⁸, might be defined as well; the generality of our approach is not affected by this choice.

Diversity based selection

We demonstrate the method for boron, one of the most structurally complex elements⁵⁹. With the exception of a high-pressure α -Ga type phase, all relevant boron allotropes contain B_{12} icosahedra as the defining structural unit⁵⁹. Boron has been the topic of structure searches with DFT^{60–63} and, more recently, with ML potentials for bulk allotropes^{18,20} and gas-phase clusters²². Our previous work showed how the PES for boron can be fitted in a ML framework¹⁸, leading to an interatomic potential able to describe the different allotropes. However, at that time, we generated and fed back 250 cells per iteration (without further selection), and added the structure of α - B_{12} manually at a later stage.¹⁸

Our new protocol “discovers” the structure of α -B₁₂ in a self-guided way, as shown in Fig. 2. The figure compares the performance of our selection procedure with alternatives: (i) random selection, and (ii) using CUR but on the matrix of SOAP vectors rather than similarity kernels (see Methods section for details). The first of these, random selection, improves the database much less after the first few iterations, and ends up with the highest error (gray in Fig. 2). The second, which uses CUR but neglects the non-linear aspects of the similarity kernel, initially performs well, but soon stops reducing the error (green). Note that this algorithm is exactly the same as the one used in potential fitting to select representative environments (in that case, even computing the complete similarity kernel matrix quickly becomes impractical). The use of CUR on the similarity kernel for selecting structures to be included in the next iteration is shown to be the most efficient (purple in Fig. 2).

The increasingly accurate description of the B₁₂ icosahedron is reflected in a gradually lowered energy error, falling below the 10 meV/atom threshold with fewer than 2,000 DFT evaluations, and below 4 meV/atom once the cycle is completed. This improvement is best understood by inspecting the respective lowest-energy structures that enter the database in a given iteration (Fig. 2). The lowest-energy structure at point **A** already contains several three-membered rings, but no B₁₂ icosahedra yet. With one more iteration, there is a sharp drop in the GAP error (from 175 to 51 meV/at.), concomitant with the first appearance of a rather distorted α -B₁₂ structure (**B**). The final database has seen several instances of the correctly ordered structure (**C**).

Learning diverse crystal structures

Our method is not restricted to a particular chemical system. To demonstrate this, we now apply it to three prototypical materials side by side: carbon, silicon, and titanium, which all exhibit multiple crystal structures.

In carbon (Fig. 3a), both the layered structure of graphite and the tetrahedral network of diamond are correctly “learned” during our iterations. For graphite, the energy error reaches a plateau after only a few hundred DFT evaluations; for diamond, the initial error is very large, and after a dozen or so iterations we observe a rapid drop—concomitant with a drop in the error for the structurally very similar lonsdaleite (“hexagonal diamond”). The

163 final prediction error is well below 1 meV/atom for the sp^3 bonded allotropes, and on the
 164 order of 4 meV/atom for graphite. We have previously shown that the forces in diamond
 165 show higher locality than those in graphite, making their description by a finite-ranged
 166 ML potential easier³², given that sufficient training data are available. We also note that
 167 our method captures the difference between diamond and lonsdaleite very well: its value is
 168 27 meV/atom with the final GAP-RSS version, and 28 meV/atom with DFT.

169 In silicon (Fig. 3b), the ground-state (diamond-type) structure is very quickly learned,
 170 more quickly so than diamond carbon, which we ascribe to the absence of a competing
 171 threefold-coordinated phase in the case of Si. We further test our evolving potentials on
 172 the high-pressure form, the β -tin type allotrope (space group $I4_1/amd$), which is easily
 173 discovered; the larger residual error for β -Sn-type than for diamond-type Si is consistent
 174 with previous studies using a manually tuned potential³⁴. We also test our method on a
 175 recently synthesized open-framework structure with 24 atoms in the unit cell (oS24)⁶⁵, which
 176 consists of distorted tetrahedral building units that are linked in different ways, which the
 177 potential has not “seen”. Still, a good description is achieved after a few iterations.

178 In titanium (Fig. 3c), a hexagonal close packed (hcp) structure is observed at ambient
 179 conditions; however, the zero-Kelvin ground state has been under debate: depending on the
 180 DFT method, either hcp or the so-called ω phase is obtained as the minimum. Our method
 181 clearly reproduces the qualitative and quantitative difference between the two allotropes
 182 (22 meV/atom with the final GAP-RSS iteration versus 24 meV/atom with DFT) at the
 183 computational level we use, viz. PBEsol⁶⁶.

184 Looking beyond the minimum structures, the DFT energy–volume curves are, by and
 185 large, well reproduced by GAP-RSS; see Fig. 3d–f. There is some deviation at large volumes
 186 for hcp and ω -type Ti, but this is an acceptable issue as these regions of the PES are not
 187 as relevant, corresponding to negative external pressure. If one were interested in very
 188 accurate elastic properties, one would choose to include less dense structures by modifying
 189 the pressure parameters (Methods section, Eq. 5). Indeed, it was recently shown that a ML
 190 potential for Ti, fitted to a database of 2,700 structures built from the phases on which we
 191 test here (ω , hcp, bcc) and other relevant structures can make an accurate prediction of
 192 energetic and elastic properties⁶⁷.

While the most relevant crystal structures for materials are usually well known and available from databases, we show that our chemically “agnostic” approach is more general. In Fig. 4, we present an energy–energy scatter plot for the last set of GAP-RSS minimizations, evaluated with DFT and with the preceding GAP version, and again across three different chemical systems. We survey both the low- and higher-energy regions of the PES—up to 1 eV per atom, which is very roughly the upper stability limit at which crystalline carbon phases may be expected to exist⁶⁸. The higher-energy regions clearly exhibit a larger error; when generating a potential for specific crystalline phases, one might choose to exclude them at a later stage. We specifically do not exclude high-energy structures, because we aim to generate potentials that will be useful for future structure searches.

To analyze and understand the outcome of these searches in structural and chemical terms, we use a dimensionality reduction technique to draw a two-dimensional structural map. Various types of SOAP-based maps have been used with success to analyze structural and chemical relationships in different materials datasets^{56,58,69}. Here, we use them to illustrate how different materials (including their allotropes as known from chemistry textbooks) are related in structural space.

To compare different materials with inherently different *absolute* bond lengths, we rescale their unit cells such that the minimum bond length in each is $r_0 = 1.0 \text{ \AA}$, inspired by approaches for topological analyses of different structures⁷⁰. We then use kernel principal component analysis with a SOAP kernel to represent the structures in a 2D plane. Figure 5 shows the resulting plot, in which we have encoded the species by symbols and the average coordination number by color. (Coordination numbers are determined by counting nearest neighbors up to $1.2 r_0$.)

The results fall within four groups, moving from the left to the right through Fig. 5. The first group is given by graphite-like structures; they are three-fold coordinated and only carbon structures (circles) are found there. Roman numerals in Fig. 5 indicate examples, and in this first group we observe flat (i) and buckled (ii) graphite sheets. In the second group, we have four-fold coordinated (“diamond-like”) networks, made up by both carbon and silicon (recall that we are using a normalized bond length, so diamond carbon and diamond-type silicon will fall on the same position in the plot). The structures that are shown as insets

are characteristic examples; from left to right, there is a distorted lonsdaleite-type structure (iii), the well-known **unj** framework (also referred to as the “chiral framework structure” in group-14 elements (iv)⁷¹, and a more complex sp^3 -bonded allotrope (v). While the axis values in our plot are arbitrary, they naturally reflect the structural evolution toward higher coordination numbers, and therefore we next observe a set of high-pressure silicon structures (squares), such as the simple-hexagonal one (vi), with an additional contribution from lower-coordinated titanium structures (circles). Finally, there is a set of densely packed structures, all clustered closely together; these are titanium structures including hcp (vii) and the ω type (viii). In the center of the plot, there is a structure that bears resemblance to none of the previously mentioned ones (ix), an energetically high-lying and strongly disordered intermediate from a relaxation trajectory that was added to the reference database, rather than a local minimum (see also Tables S1–3). This dissimilarity is reflected in relatively large distances from other entries in the SOAP-based similarity map.

DISCUSSION

We have shown that automated protocols can be designed for generating structural databases and fitting potential-energy surfaces of materials in a self-guided way. This allows for the generation of ML-based interatomic potentials with minimal effort, both in terms of computational and user time, when combined with a suitable fitting framework, of which many are presently available. Formalizing the protocols for database construction is an important step toward further methodological developments and, ultimately, toward wide applicability of these techniques in computational materials science.

Our RSS-based reference databases efficiently cover structural space up to a given system size (here, 24 atoms in the unit cell). Once a core database has been constructed in this way, it may be readily improved by adding defect, surface, and liquid/amorphous structural models in much larger simulation cells, while at the same time being sufficiently robust to avoid unphysical behavior—even when taken to the more extreme regions of configuration space that are explored early on during RSS.

We targeted here the space of three-dimensional inorganic crystal structures, but conceptually similar approaches may be useful for nanoparticles^{23,72} and other lower-dimensional systems. Finally, organic (molecular) materials are also beginning to be described very re-

liably with ML potentials^{7,11}, and an interesting open question is how to use the structural diversity inherent in RSS in the context of organic solids⁷³.

METHODS

Interatomic potential fitting

To fit interatomic potentials, we use the established Gaussian Approximation Potential (GAP) ML framework⁴ and the associated computer code, which is freely available for non-commercial research at <http://www.libatoms.org>. Compared to previous work, we here use suitable heuristics to automate and generalize the choice of fitting parameters where possible. We stress again, however, that the main development in the present work is in the automated generation of databases, not the descriptors or the regressor.

We use a linear combination of 2-, 3-, and many-body terms following Refs. 74 and 32, with defining parameters given in Table 1. The 2-body (“2b”) and 3b descriptors are scalar distances and symmetrized three-component vectors, respectively. For the many-body term, we use the Smooth Overlap of Atomic Positions (SOAP) kernel⁵⁵, which has been used to fit GAPs for diverse systems^{28,32–34}. The overall energy scale of each descriptor’s contribution to the predicted energy (controlled by the parameter δ)⁷⁴ is set automatically in our protocol. The 2b value is set from the variance of energies in the fitting database, the 3b value is set from the energy error between a 2b-only fit and the fitting database, and the SOAP value is set from the energy error for a 2b+3b-only fit.

The cutoffs for the three types of descriptors are expressed in terms of the characteristic radius r (Table 1): that for 2b is longest range, while that for 3b is shortest (intended to capture only nearest neighbors), and the SOAP is intermediate in range. The resulting cutoff settings are listed in Table 1, the characteristic radii r for the systems studied here being 0.84, 0.76, 1.11, and 1.47 Å for B, C, Si, and Ti, respectively. An ad-hoc choice is made here between predominantly covalent (B, C, Si) or metallic (Ti) materials for selecting the appropriate tabulated radii; however, settings based on the covalent radius for silicon also produce a satisfactory fit for the metallic (β -tin type) modification (residual error < 10 meV / at.; Fig. 3b). Future work might explore more automated ways of extracting optimal atomic radii from datasets, and suitable definitions

for multicomponent systems (we stress that the latter, in principle, can be routinely treated by present-day ML potentials^{14,37,46}). None of this is expected to affect the conclusions of the present work.

The weights on the energies, forces, and stresses that are fit are set by diagonal noise terms in Gaussian process regression⁴. We set these according to the reference energy of a given structure, to make the fit more accurate for relatively low-energy structures at each volume while providing flexibility for the higher-energy regions. The values are piecewise-linear functions in ΔE , which is the per-atom reference energy difference relative to the same volume on the convex hull bounding the set of (V, E) points from below (in energy). For the energy the error σ_E is 1 meV/atom for $\Delta E \leq 0.1$ eV, 100 meV/atom for $\Delta E \geq 1$ eV, and linearly interpolated in between. For forces the corresponding σ_F values are 31.6 and 316 meV/Å, and for virials the σ_V values are 63.2 and 632 meV/atom.

Comparing structures

The same mathematical tools that are used to compare atomic environments for the purpose of constructing potentials can also be used to compare atomic configurations⁵⁶. As for the regression, for these similarity kernels we also use SOAP, although with different parameters ($n_{\max} = l_{\max} = 12$, $\sigma_{\text{at}} = 0.0875$ Å, $r_{\text{cut}} = 10.5$ Å), to compare the similarity of environments in selecting from which data to train (in the CUR step). For the kernel PCA used to generate the map in Fig. 5, we use $n_{\max} = l_{\max} = 16$, $\sigma = 0.1$ r_0 and $r_{\text{cut}} = 2.5$ r_0 , where r_0 is the shortest bond length, as described in the Results section. We obtain what we call a “configuration-averaged” SOAP by averaging over all atoms in the cell. In the SOAP framework⁵⁵, the neighbor density of a given atom i is expanded using a local basis set of radial basis functions g_n and spherical harmonics Y_{lm} ,

$$\begin{aligned}\rho_i(\mathbf{r}) &= \sum_j \exp(-|\mathbf{r} - \mathbf{r}_{ij}|^2/2\sigma_{\text{at}}^2) \\ &= \sum_{nlm} c_{nlm}^{(i)} g_n(r) Y_{lm}(\hat{\mathbf{r}}),\end{aligned}\tag{1}$$

where j runs over the neighbours of atom i within the specified cutoff (including i itself). To obtain a similarity measure between unit cells, rather than individual atoms, we then average the

expansion coefficients over all atoms a in the unit cell,

$$\bar{c}_{nlm} = \frac{1}{N} \sqrt{\frac{8\pi^2}{2l+1}} \sum_i c_{nlm}^{(i)}, \quad (2)$$

and construct the rotationally invariant power spectrum for the entire unit cell⁵⁷,

$$\bar{p}_{nn'l} = \sum_m (\bar{c}_{nlm})^* \bar{c}_{n'lm}. \quad (3)$$

Note that this is not equal to the average of the usual atomic SOAP power spectra used to describe the atomic neighbor environments. The final kernel to compare two cells, A and B, is then

$$k_{AB} = \left(\sum_{nn'l} \bar{p}_{nn'l}^{(\text{cell A})} \bar{p}_{nn'l}^{(\text{cell B})} \right)^\zeta, \quad (4)$$

where ζ is a small integer number (here, $\zeta = 4$).

For our main results, our diverse structure selection uses leverage-score CUR⁵² applied to the matrix of similarity kernels between atomic configurations. We also test a version of our method where the CUR algorithm is applied to the rectangular matrix of configuration-averaged SOAP *vectors*, rather than the square matrix of similarity kernels. This qualitatively captures the same information, but neglects the non-linear nature of the exponentiation that transforms the (linear) dot-product of SOAP vectors into the similarity kernel. The results of these methods are compared in Figure 2 and Figures S1–2.

Iterative generation of reference data

Randomized atomic positions are generated using the `buildcell` code of the AIRSS package version 0.9, available at <https://www.mtg.msm.cam.ac.uk/Codes/AIRSS>. The positions are repeated by 1–8 symmetry operations, and the cells contain 6–24 atoms. A minimum separation is also set, with a value of $1.8r$. The volumes per atom of the random cells are centered on $V_0 = 14.5r^3$ for covalent, and $V_0 = 5.5r^3$ for metallic systems. In the initial iteration, half of the structures are generated from the `buildcell`-default narrow range of volumes, and half from a wider range, $\pm 25\%$ from the heuristic value. In all later iterations, only the default narrow range is used. The wide volume range configurations are meant to simply span a wide range of structures¹⁸, and use only even numbers of atoms. The narrow volume range configurations are meant to be good initial

conditions for RSS, and so for 80% (20%) of the seed structures, we choose even (odd) numbers of atoms, respectively. This is because for most known structures, the number of atoms in the conventional unit cell is even (eight for diamond and rocksalt, for example), although for some it is odd, including the ω phase⁷⁵. Biasing initial seeds toward distributions that occur in nature is a central idea within the AIRSS formalism⁴². The setup of these cells, in itself, has negligible computational cost compared to the relaxations: generating 10,000 candidate structures required less than 5 minutes on 16 cores (and constructing the SOAP vectors for structural selection required on the order of one minute). For the computational cost of potential fitting, see Figure S3.

With the initial potential available, we then run structural optimizations by relaxing the candidate configurations with a preconditioned LBFGS algorithm⁷⁶ to minimize the enthalpy until residual forces fall below 0.01 eV/Å. As in Ref. 19, we employ a random external pressure p with probability density

$$P(p/p_0) = \frac{1}{\beta} \exp\left(-\frac{1}{\beta} p/p_0\right), \quad (5)$$

here with $p_0 = 1$ GPa, and $\beta = 0.2$. This protocol ensures that there is a small but finite external pressure, and also some smaller-volume structures are included in the fit^{18,19}. We choose the same pressure range for all materials, for simplicity, although this value could be adjusted depending on the pressure region of interest.

The selection of configurations for DFT evaluation and fitting at each iteration involves a Boltzmann-biased flat histogram and leverage-score CUR, as illustrated in Fig. 1. To compute the selection probabilities for the flat histogram stage, the distribution of enthalpies (each computed using the pressure at which the corresponding RSS minimization was done) is approximated by the numpy⁷⁷ histogram function, with default parameters. The probability of selecting each configuration is inversely proportional to the density of the corresponding histogram bin, multiplied by a Boltzmann biasing factor. The biasing factor is exponential in the enthalpy per atom relative to the lowest enthalpy configuration, divided by a temperature of 0.3 eV for the first iteration, 0.2 eV for the second, and 0.1 eV for all remaining iterations. The leverage-score CUR selection is based on the singular-value decomposition of the square kernel matrix using the SOAP descriptors (with the dot-product kernel and exponentiation by ζ , Eq. 4). Applying the same algorithm to the rectangular matrix of SOAP descriptor vectors was significantly less effective (Figure 2).

Computational details

Reference energies and forces were obtained using DFT, with projector augmented-waves (PAW)^{78,79} as implemented in the Vienna Ab Initio Simulation Package (VASP)⁸⁰. Valence electrons were described by plane-wave basis sets with cutoff energies of 500 (B), 800 (C), 400 (Si), and 285 eV (Ti), respectively. Reciprocal space was sampled using a fixed “KSPACING” parameter in VASP, amounting to 0.25 for B, Si, and Ti, and 0.35 for C (in units of \AA^{-1} along the reciprocal lattice vectors which include the 2π factor). Exchange and correlation were treated using the PBEsol functional⁶⁶ for all materials except carbon, where the opt-B88-vdW functional^{81–83} was chosen to properly account for the van der Waals interactions in graphitic structures. Benchmark data for energy–volume curves were obtained by scaling selected unit cells within given volume increments and optimizing while constraining the volume and symmetry of the cell.

DATA AVAILABILITY

Data supporting this publication are available at <https://doi.org/10.17863/CAM.43407>.

CODE AVAILABILITY

A Python implementation for the protocol developed in this publication is available at <https://doi.org/10.17863/CAM.43407>.

ACKNOWLEDGMENTS

We thank Profs. C. J. Pickard and D. M. Proserpio for ongoing valuable discussions. N.B. acknowledges support from the Office of Naval Research through the U. S. Naval Research Laboratory’s core basic research program. G.C. acknowledges EPSRC grant EP/P022596/1. V.L.D.

acknowledges a Leverhulme Early Career Fellowship and support from the Isaac Newton Trust.

COMPETING INTERESTS

The authors declare no competing financial or non-financial interests.

AUTHOR CONTRIBUTIONS

N.B., G.C., and V.L.D. jointly designed the research, developed the approach, and analyzed the data. N.B. developed the computational framework and performed the computations with input from all authors. V.L.D. wrote the paper with input from all authors. All authors revised the paper and approved its final version.

* vld24@cam.ac.uk; *Present address:* Department of Chemistry, University of Oxford, Oxford OX1 3QR, United Kingdom

¹ Lejaeghere, K. et al. Reproducibility in density functional theory calculations of solids. *Science* **351**, aad3000 (2016).

² Oganov, A. R.; Pickard, C. J.; Zhu, Q. & Needs, R. J. Structure prediction drives materials discovery. *Nat. Rev. Mater.* **4**, 331–348 (2019).

³ Behler, J. & Parrinello, M. Generalized neural-network representation of high-dimensional potential-energy surfaces. *Phys. Rev. Lett.* **98**, 146401 (2007).

⁴ 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).

⁵ Li, Z.; Kermode, J. R. & De Vita, A. Molecular dynamics with on-the-fly machine learning of quantum-mechanical forces. *Phys. Rev. Lett.* **114**, 096405 (2015).

⁶ Artrith, N. & Urban, A. An implementation of artificial neural-network potentials for atomistic materials simulations: Performance for TiO₂. *Comput. Mater. Sci.* **114**, 135–150 (2016).

⁷ Smith, J. S.; Isayev, O. & Roitberg, A. E. ANI-1: an extensible neural network potential with DFT accuracy at force field computational cost. *Chem. Sci.* **8**, 3192–3203 (2017).

- 392 ⁸ Chmiela, S. et al. Machine learning of accurate energy-conserving molecular force fields. *Sci.*
393 *Adv.* **3**, e1603015 (2017).
- 394 ⁹ Behler, J. First principles neural network potentials for reactive simulations of large molecular
395 and condensed systems. *Angew. Chem. Int. Ed.* **56**, 12828–12840 (2017).
- 396 ¹⁰ Huan, T. D. A universal strategy for the creation of machine learning-based atomistic force
397 fields. *npj Comput. Mater.* **3**, 37 (2017).
- 398 ¹¹ Chmiela, S.; Sauceda, H. E.; Müller, K.-R. & Tkatchenko, A. Towards exact molecular dynamics
399 simulations with machine-learned force fields. *Nat. Commun.* **9**, 3887 (2018).
- 400 ¹² Zhang, L.; Han, J.; Wang, H.; Car, R. & E, W. Deep potential molecular dynamics: A scalable
401 model with the accuracy of quantum mechanics. *Phys. Rev. Lett.* **120**, 143001 (2018).
- 402 ¹³ Caro, M. A.; Deringer, V. L.; Koskinen, J.; Laurila, T. & Csányi, G. Growth mechanism and
403 origin of high sp^3 content in tetrahedral amorphous carbon. *Phys. Rev. Lett.*, **120**, 166101
404 (2018).
- 405 ¹⁴ Hellström, M.; Quaranta, V. & Behler, J. One-dimensional vs. two-dimensional proton transport
406 processes at solid–liquid zinc-oxide–water interfaces. *Chem. Sci.* **10**, 1232–1243 (2019).
- 407 ¹⁵ Feller, M. R.; Tan, A. M. Z.; Hector, L. G. & Trinkle, D. R. Geometries of edge and mixed
408 dislocations in bcc Fe from first-principles calculations. *Phys. Rev. Mater.* **2**, 113605 (2018).
- 409 ¹⁶ Maresca, F.; Dragoni, D.; Csányi, G.; Marzari, N. & Curtin, W. A. Screw dislocation structure
410 and mobility in body centered cubic Fe predicted by a gaussian approximation potential. *npj*
411 *Comput. Mater.* **4**, 69 (2018).
- 412 ¹⁷ Deringer, V. L.; Csányi, G. & Proserpio, D. M. Extracting crystal chemistry from amorphous
413 carbon structures. *ChemPhysChem* **18**, 873–877 (2017).
- 414 ¹⁸ Deringer, V. L.; Pickard, C. J. & Csányi, G. Data-driven learning of total and local energies in
415 elemental boron. *Phys. Rev. Lett.* **120**, 156001 (2018).
- 416 ¹⁹ Deringer, V. L.; Proserpio, D. M.; Csányi, G. & Pickard, C. J. Data-driven learning and pre-
417 diction of inorganic crystal structures. *Faraday Discuss.* **211**, 45–59 (2018).
- 418 ²⁰ Podryabinkin, E. V.; Tikhonov, E. V.; Shapeev, A. V. & Oganov, A. R. Accelerating crystal
419 structure prediction by machine-learning interatomic potentials with active learning. *Phys. Rev.*
420 *B* **99**, 064114 (2019).
- 421 ²¹ Ouyang, R.; Xie, Y. & Jiang, D.-e. Global minimization of gold clusters by combining neural
422 network potentials and the basin-hopping method. *Nanoscale* **7**, 14817–14821 (2015).

- 423 ²² Tong, Q.; Xue, L.; Lv, J.; Wang, Y. & Ma, Y. Accelerating CALYPSO structure prediction by
424 data-driven learning of a potential energy surface. *Faraday Discuss.* **211**, 31–43 (2018).
- 425 ²³ Kolsbjerg, E. L.; Peterson, A. A. & Hammer, B. Neural-network-enhanced evolutionary algo-
426 rithm applied to supported metal nanoparticles. *Phys. Rev. B* **97**, 195424 (2018).
- 427 ²⁴ Hajinazar, S.; Sandoval, E. D.; Cullo, A. J. & Kolmogorov, A. N. Multitribe evolutionary search
428 for stable Cu–Pd–Ag nanoparticles using neural network models. *Phys. Chem. Chem. Phys.* **21**,
429 8729–8742 (2019).
- 430 ²⁵ Eivari, H. A. et al. Two-dimensional hexagonal sheet of TiO₂. *Chem. Mater.* **29**, 8594–8603
431 (2017).
- 432 ²⁶ Behler, J.; Martoňák, R.; Donadio, D. & Parrinello, M. Metadynamics simulations of the high-
433 pressure phases of silicon employing a high-dimensional neural network potential. *Phys. Rev.*
434 *Lett.* **100**, 185501 (2008).
- 435 ²⁷ Behler, J.; Martoňák, R.; Donadio, D. & Parrinello, M. Pressure-induced phase transitions in
436 silicon studied by neural network-based metadynamics simulations. *Phys. Status Solidi B* **245**,
437 2618–2629 (2008).
- 438 ²⁸ Szlachta, W. J.; Bartók, A. P. & Csányi, G. Accuracy and transferability of Gaussian approxi-
439 mation potential models for tungsten. *Phys. Rev. B* **90**, 104108 (2014).
- 440 ²⁹ Dragoni, D.; Daff, T. D.; Csányi, G. & Marzari, N. Achieving DFT accuracy with a machine-
441 learning interatomic potential: Thermomechanics and defects in bcc ferromagnetic iron. *Phys.*
442 *Rev. Mater.* **2**, 013808 (2018).
- 443 ³⁰ Eshet, H.; Khaliullin, R. Z.; Kühne, T. D.; Behler, J. & Parrinello, M. Ab initio quality neural-
444 network potential for sodium. *Phys. Rev. B* **81**, 184107 (2010).
- 445 ³¹ Sosso, G. C.; Miceli, G.; Caravati, S.; Behler, J. & Bernasconi, M. Neural network interatomic
446 potential for the phase change material GeTe. *Phys. Rev. B* **85**, 174103 (2012).
- 447 ³² Deringer, V. L. & Csányi, G. Machine learning based interatomic potential for amorphous
448 carbon. *Phys. Rev. B* **95**, 094203 (2017).
- 449 ³³ Mocanu, F. C. et al. Modeling the phase-change memory material, Ge₂Sb₂Te₅, with a machine-
450 learned interatomic potential. *J. Phys. Chem. B* **122**, 8998–9006 (2018).
- 451 ³⁴ Bartók, A. P.; Kermode, J.; Bernstein, N. & Csányi, G. Machine learning a general-purpose
452 interatomic potential for silicon. *Phys. Rev. X* **8**, 041048 (2018).

453 ³⁵ Bartók, A. P. et al. Machine learning unifies the modeling of materials and molecules. *Sci. Adv.*
454 **3**, e1701816 (2017).

455 ³⁶ Deringer, V. L. et al. Realistic atomistic structure of amorphous silicon from machine-learning-
456 driven molecular dynamics. *J. Phys. Chem. Lett.* **9**, 2879–2885 (2018).

457 ³⁷ Hajinazar, S.; Shao, J. & Kolmogorov, A. N. Stratified construction of neural network based
458 interatomic models for multicomponent materials. *Phys. Rev. B* **95**, 014114 (2017).

459 ³⁸ Jinnouchi, R.; Lahnsteiner, J.; Karsai, F.; Kresse, G. & Bokdam, M. Phase transitions of
460 hybrid perovskites simulated by machine-learning force fields trained on-the-fly with Bayesian
461 inference. *Phys. Rev. Lett.* **122**, 225701 (2019).

462 ³⁹ Vandermause, J. & Torrisi, S. B.; Batzner, S.; Kolpak, A. M. & Kozinsky, B. On-the-fly
463 Bayesian active learning of interpretable force-fields for atomistic rare events. Preprint at
464 <https://arxiv.org/abs/1904.02042> (2019).

465 ⁴⁰ Jinnouchi, R.; Karsai, F. & Kresse, G. On-the-fly machine learning force field generation: Ap-
466 plication to melting points. *Phys. Rev. B* **100**, 014105 (2019).

467 ⁴¹ Pickard, C. J. & Needs, R. J. High-pressure phases of silane. *Phys. Rev. Lett.* **97**, 045504 (2006).

468 ⁴² Pickard, C. J. & Needs, R. J. Ab initio random structure searching. *J. Phys.: Condens. Matter*
469 **23**, 053201 (2011).

470 ⁴³ Artrith, N. & Behler, J. High-dimensional neural network potentials for metal surfaces: A
471 prototype study for copper. *Phys. Rev. B* **85**, 045439 (2012).

472 ⁴⁴ Shapeev, A. Moment tensor potentials: A class of systematically improvable interatomic poten-
473 tials. *Multiscale Model. Simul.* **14**, 1153–1173 (2016).

474 ⁴⁵ Podryabinkin, E. V. & Shapeev, A. V. Active learning of linearly parametrized interatomic
475 potentials. *Comput. Mater. Sci.* **140**, 171–180 (2017).

476 ⁴⁶ Gubaev, K.; Podryabinkin, E. V.; Hart, G. L. W. & Shapeev, A. V. Accelerating high-
477 throughput searches for new alloys with active learning of interatomic potentials. *Comput.*
478 *Mater. Sci.* **156**, 148–156 (2019).

479 ⁴⁷ Zhang, L.; Lin, D.-Y.; Wang, H.; Car, R. & E, W. Active learning of uniformly accurate
480 interatomic potentials for materials simulation. *Phys. Rev. Mater.* **3**, 023804 (2019).

481 ⁴⁸ Cordero, B. et al. Covalent radii revisited. *Dalton Trans.*, 2832–2838 (2008).

482 ⁴⁹ Pickard, C. J. & Needs, R. Highly compressed ammonia forms an ionic crystal. *Nat. Mater.* **7**,
483 775–779 (2008).

484 ⁵⁰ Marqués, M. et al. Crystal structures of dense lithium: A metal-semiconductor-metal transition.
485 *Phys. Rev. Lett.* **106**, 095502 (2011).

486 ⁵¹ Stratford, J. M. et al. Investigating sodium storage mechanisms in tin anodes: A combined pair
487 distribution function analysis, density functional theory, and solid-state NMR approach. *J. Am.*
488 *Chem. Soc.* **139**, 7273–7286 (2017).

489 ⁵² Mahoney, M. W. & Drineas, P. CUR matrix decompositions for improved data analysis. *Proc.*
490 *Natl. Acad. Sci., U. S. A.* **106**, 697–702 (2009).

491 ⁵³ Mones, L.; Bernstein, N. & Csányi, G. Exploration, sampling, and reconstruction of free energy
492 surfaces with Gaussian process regression. *J. Chem. Theory Comput.* **12**, 5100–5110 (2016).

493 ⁵⁴ Imbalzano, G. et al. Automatic selection of atomic fingerprints and reference configurations for
494 machine-learning potentials. *J. Chem. Phys.* **148**, 241730 (2018).

495 ⁵⁵ Bartók, A. P.; Kondor, R. & Csányi, G. On representing chemical environments. *Phys. Rev. B*
496 **87**, 184115 (2013).

497 ⁵⁶ De, S.; Bartók, A. P.; Csányi, G. & Ceriotti, M. *Phys. Chem. Chem. Phys.* **18**, 13754–13769
498 (2016).

499 ⁵⁷ Mavračić, J.; Mocanu, F. C.; Deringer, V. L.; Csányi, G. & Elliott, S. R. Similarity between
500 amorphous and crystalline phases: The case of TiO₂. *J. Phys. Chem. Lett.* **9**, 2985–2990 (2018).

501 ⁵⁸ Caro, M. A.; Aarva, A.; Deringer, V. L.; Csányi, G. & Laurila, T. Reactivity of amorphous
502 carbon surfaces: Rationalizing the role of structural motifs in functionalization using machine
503 learning. *Chem. Mater.* **30**, 7446–7455 (2018).

504 ⁵⁹ Albert, B. & Hillebrecht, H. Boron: Elementary challenge for experimenters and theoreticians.
505 *Angew. Chem. Int. Ed.* **48**, 8640–8668 (2009).

506 ⁶⁰ Oganov, A. R. et al. Ionic high-pressure form of elemental boron. *Nature* **457**, 863–867 (2009).

507 ⁶¹ Wu, X. et al. Two-dimensional boron monolayer sheets. *ACS Nano* **6**, 7443–7453 (2012).

508 ⁶² Mannix, A. J. et al. Synthesis of borophenes: Anisotropic, two-dimensional boron polymorphs.
509 *Science* **350**, 1513–1516 (2015).

510 ⁶³ Ahnert, S. E.; Grant, W. P.; Pickard, C. J. Revealing and exploiting hierarchical material
511 structure through complex atomic networks. *npj Comput. Mater.* **3**, 35 (2017).

512 ⁶⁴ Momma, K. & Izumi, F. VESTA 3 for three-dimensional visualization of crystal, volumetric
513 and morphology data. *J. Appl. Crystallogr.* **44**, 1272–1276 (2011).

514 ⁶⁵ Kim, D. Y.; Stefanoski, S.; Kurakevych, O. O. & Strobel, T. A. Synthesis of an open-framework
515 allotrope of silicon. *Nat. Mater.* **14**, 169–173 (2015).

516 ⁶⁶ Perdew, J. P. et al. Restoring the density-gradient expansion for exchange in solids and surfaces.
517 *Phys. Rev. Lett.* **100**, 136406 (2008).

518 ⁶⁷ Takahashi, A.; Seko, A. & Tanaka, I. Conceptual and practical bases for the high accuracy of
519 machine learning interatomic potentials: Application to elemental titanium. *Phys. Rev. Mater.*
520 **1**, 063801 (2017).

521 ⁶⁸ Aykol, M.; Dwaraknath, S. S.; Sun, W. & Persson, K. A. Thermodynamic limit for synthesis of
522 metastable inorganic materials. *Sci. Adv.* **4**, eaaq0148 (2018).

523 ⁶⁹ Engel, E. A.; Anelli, A.; Ceriotti, M.; Pickard, C. J. & Needs, R. J. Mapping uncharted territory
524 in ice from zeolite networks to ice structures. *Nat. Commun.* **9**, 2173 (2018).

525 ⁷⁰ Delgado-Friedrichs, O. & O’Keeffe, M. Identification of and symmetry computation for crystal
526 nets. *Acta Crystallogr., Sect. A* **59**, 351–360 (2003).

527 ⁷¹ Pickard, C. J. & Needs, R. J. Hypothetical low-energy chiral framework structure of group 14
528 elements. *Phys. Rev. B* **81**, 014106 (2010).

529 ⁷² Artrith, N. & Kolpak, A. M. Understanding the composition and activity of electrocatalytic
530 nanoalloys in aqueous solvents: A combination of DFT and accurate neural network potentials.
531 *Nano Lett.* **14**, 2670–2676 (2014).

532 ⁷³ Zilka, M. et al. Ab initio random structure searching of organic molecular solids: assessment
533 and validation against experimental data. *Phys. Chem. Chem. Phys.* **19**, 25949–25960 (2017).

534 ⁷⁴ Bartók, A. P. & Csányi, G. Gaussian approximation potentials: A brief tutorial introduction.
535 *Int. J. Quantum Chem.* **115**, 1051–1057 (2015).

536 ⁷⁵ Sikka, S. K.; Vohra, Y. K. & Chidambaram, R. Omega phase in materials. *Prog. Mater. Sci.*
537 **27**, 245–310 (1982).

538 ⁷⁶ Packwood, D. et al. A universal preconditioner for simulating condensed phase materials. *J.*
539 *Chem. Phys.* **144**, 164109 (2016).

540 ⁷⁷ The numpy python library version 1.15.2, <http://www.numpy.org>.

541 ⁷⁸ Blöchl, P. E. Projector augmented-wave method. *Phys. Rev. B* **50**, 17953–17979 (1994).

542 ⁷⁹ Kresse, G. & Joubert, D. From ultrasoft pseudopotentials to the projector augmented-wave
543 method. *Phys. Rev. B* **59**, 1758–1775 (1999).

544 ⁸⁰ Kresse, G. & Furthmüller, J. Efficient iterative schemes for ab initio total-energy calculations
 545 using a plane-wave basis set. *Phys. Rev. B* **54**, 11169–11186 (1996).
 546 ⁸¹ Dion, M.; Rydberg, H.; Schröder, E.; Langreth, D. C. & Lundqvist, B. I. Van der Waals density
 547 functional for general geometries. *Phys. Rev. Lett.* **92**, 246401 (2004).
 548 ⁸² Román-Pérez, G. & Soler, J. M. Efficient implementation of a van der Waals density functional:
 549 Application to double-wall carbon nanotubes. *Phys. Rev. Lett.* **103**, 096102 (2009).
 550 ⁸³ Klimeš, J. ; Bowler, D. R. & Michaelides, A. Van der Waals density functionals applied to solids.
 551 *Phys. Rev. B* **83**, 195131 (2011).

552 **FIGURE LEGENDS**

553 **A. Figure 1**

554 An automated protocol that iteratively explores structural space and fits machine learn-
 555 ing (ML) based interatomic potentials. (a) General overview of the approach. From an
 556 ensemble of randomized unit cells (*blue*), we select the most geometrically diverse ones us-
 557 ing the leverage-score CUR algorithm. Selected cells are evaluated with single-point DFT
 558 computations and used to fit an initial Gaussian Approximation Potential (GAP) (*orange*).
 559 Then, this potential is used to relax a new ensemble of randomized cells (*green*), selecting
 560 again the most relevant snapshots, and repeating the cycle. (b) Illustration of the multi-step
 561 selection procedure. We first consider all trajectories in a given generation, sketched by con-
 562 nected points, and select local minima (using an energy criterion, the flattened histogram,
 563 and then a structural criterion, the CUR). From the trajectories leading to these minima,
 564 we then select the most representative cells; these can be intermediates (*green*) or end points
 565 (*purple*) of relaxations. The structures finally selected (*magenta*) are DFT-evaluated and
 566 added to the database.

567 **Figure 2**

568 “Learning” the crystal structure of α -rhombohedral boron. *Top*: Error of iteratively
 569 generated GAP-RSS models, for the energy of the optimised ground-state structure of α -
 570 B₁₂, referenced to DFT. Three independent runs are compared: random selection of points

571 (gray), our two-step selection procedure with CUR on SOAP *vectors* (green), and the same
572 two-step procedure but with CUR on SOAP *similarity kernels* (purple). *Bottom:* Evolution
573 of the B₁₂ icosahedron as the defining structural fragment. For three points of the $N = 100$
574 cycles, having completed 400 (“A”), 500 (“B”), and 2,500 (“C”) DFT evaluations in total,
575 the respective lowest-energy structure (at the DFT level) from this iteration is shown, as
576 visualized using VESTA⁶⁴. Bonds between atoms are drawn using a cut-off of 1.9 Å; note
577 that there are further connections between the B₁₂ icosahedra with slightly larger B···B
578 distances.

579 **Figure 3**


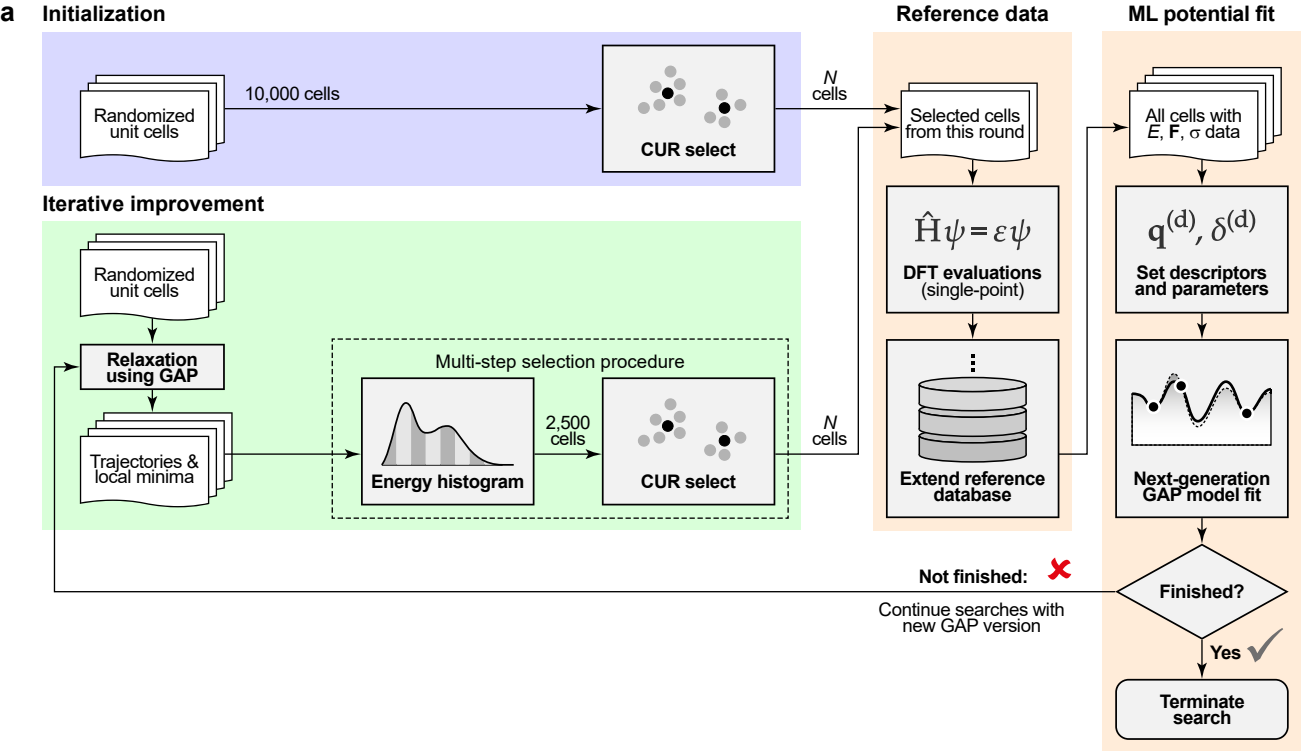
580 “Learning” diverse crystal structures without prior knowledge, including textbook ex-
581 amples of an insulator (carbon), a semiconductor (silicon), and a metal (titanium). (a–c)
582 Energy error, defined as the difference between DFT- and GAP-computed energies for struc-
583 tures optimized with the respective method. GAP-RSS models that deviate from the DFT
584 result by less than 1 meV / atom are considered to be fully converged and therefore their
585 errors are drawn as a constant minimum value to ease visualization. (d–f) Energy–volume
586 curves computed with the final GAP-RSS model (solid lines) and the DFT reference method
587 (dashed lines). The open-framework oS24 structure, at high pressure, collapses into a more
588 densely packed phase (“*”; see SI for details). All energies are referenced to the DFT result
589 for the respective most stable crystal structure.


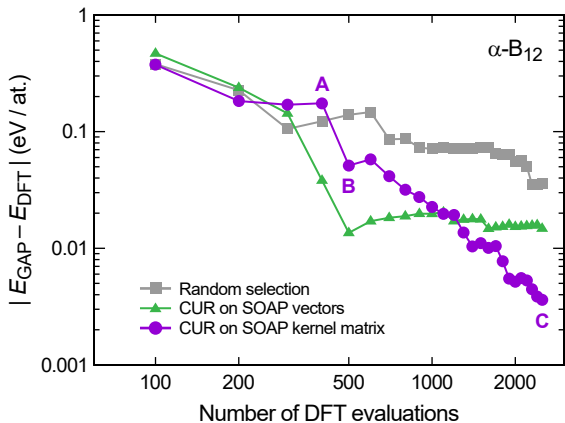
590 **Figure 4**

591 Scatter plots of predicted versus DFT energies for ensembles of structures added to the
592 reference databases in the final iteration. Note that the energy scale is continuous, but it
593 changes from linear to logarithmic scaling at 0.1 eV/at., allowing us to visualize both low-
594 and higher-energy regions.




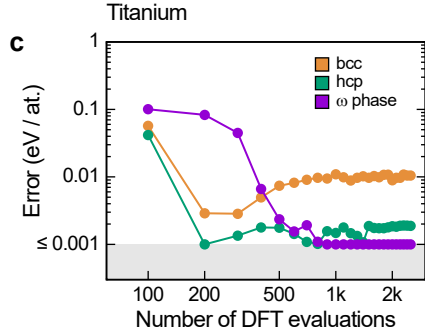
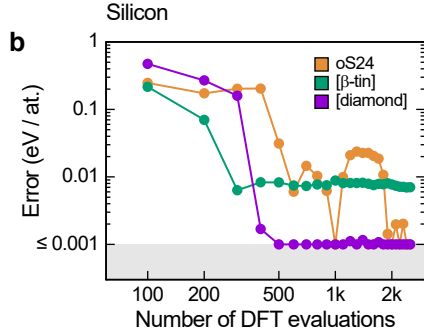
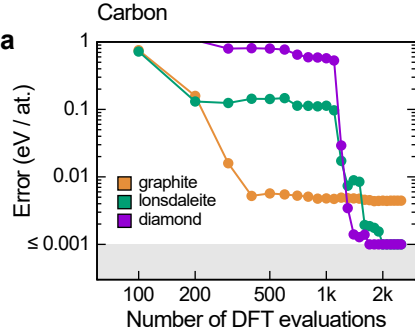
Figure 5


Visualizing the highly diverse structures, both at low and relatively high energies above the global minimum, that have been explored by GAP-RSS and added to the reference database in the last iteration. A similarity map compares three systems side-by-side (carbon, *triangles*; silicon, *squares*; titanium, *circles*), as described in the text. The resulting plot (with arbitrary axis values) emphasizes relationships between the different databases. The structures, “discovered” from scratch by our protocol, range all the way from threefold-coordinated graphite, fourfold-coordinated (sp^3 -like) allotropes of C and Si, onward to high-pressure Si structures and finally densely packed variants of Ti. A higher-energy structure (≈ 0.6 eV/at. above diamond-type silicon) from an earlier step in a minimization trajectory is included as an example, as enclosed by a dashed line.





progress of GAP-RSS iterations





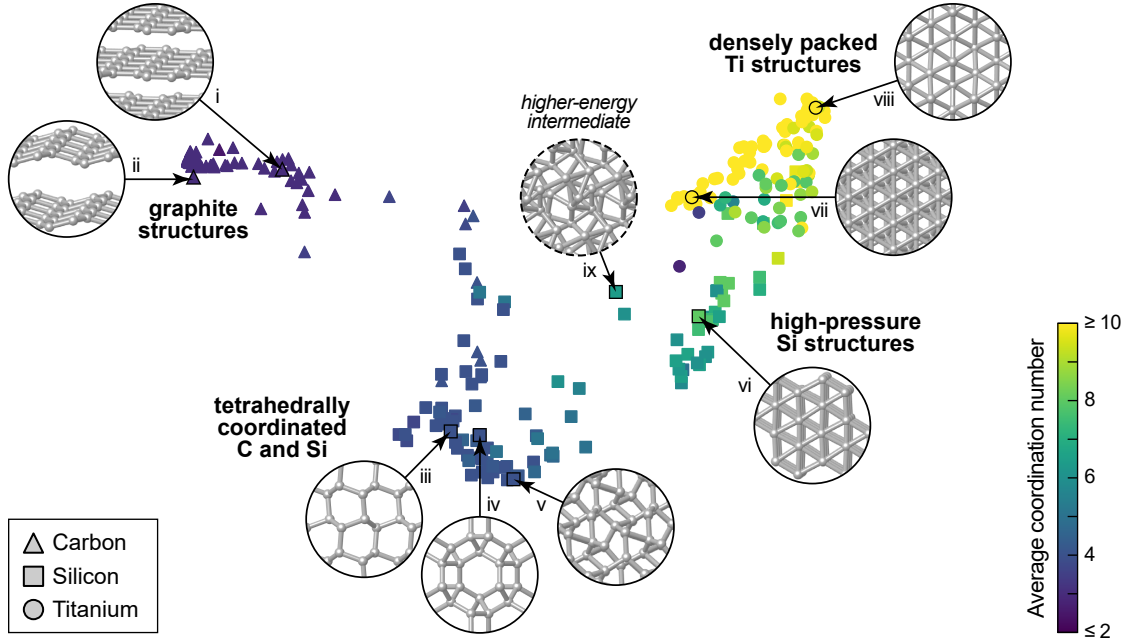


Table 1: Hyperparameters for descriptors that we use in GAP fitting. For all descriptors: Gaussian width σ_{at} (squared-exponential kernel for 2- and 3-body; atomic density width for SOAP); number of sparse points N_{sp} . For SOAP only: number of radial functions n_{max} and angular momenta l_{max} , and kernel exponent ζ . Cutoffs r_{cut} are expressed in terms of the characteristic radius r , listed for each material in the Interatomic potential fitting subsection

	σ_{at} (Å)	N_{sp}	n_{max}	l_{max}	ζ	r_{cut} (Å)	
						(covalent)	(metallic)
2-body	0.5	30				$9.0\,r$	$8.2\,r$
3-body	1.0	100				$2.925\,r$	$2.665\,r$
SOAP	0.75	2000	8	8	4	$4.5\,r$	$4.1\,r$