

## Modeling high-entropy transition metal alloys with alchemical compression

Nataliya Lopanitsyna <sup>1</sup>, Guillaume Fraux <sup>1</sup>, Maximilian A. Springer <sup>2</sup>, Sandip De <sup>2</sup>, and Michele Ceriotti <sup>1</sup>

<sup>1</sup>Laboratory of Computational Science and Modeling, Institute of Materials, École Polytechnique Fédérale de Lausanne, 1015 Lausanne, Switzerland

<sup>2</sup>BASF SE, Carl-Bosch-Straße 38, 67056 Ludwigshafen, Germany



(Received 20 December 2022; revised 3 April 2023; accepted 6 April 2023; published 26 April 2023)

Alloys composed of several elements in roughly equimolar composition, often referred to as high-entropy alloys, have long been of interest for their thermodynamics and peculiar mechanical properties, and more recently for their potential application in catalysis. They are a considerable challenge to traditional atomistic modeling, and also to data-driven potentials that for the most part have memory footprint, computational effort, and data requirements which scale poorly with the number of elements included. We apply a recently proposed scheme to compress chemical information in a lower-dimensional space, which reduces dramatically the cost of the model with negligible loss of accuracy, to build a potential that can describe 25 *d*-block transition metals. The model shows semiquantitative accuracy for prototypical alloys and is remarkably stable when extrapolating to structures outside its training set. We use this framework to study element segregation in a computational experiment that simulates an equimolar alloy of all 25 elements, mimicking the seminal experiments in the groups of Yeh and Cantor, and use our observations on the short-range order relations between the elements to define a data-driven set of Hume-Rothery rules that can serve as guidance for alloy design. We conclude with a study of three prototypical alloys, CoCrFeMnNi, CoCrFeMoNi, and IrPdPtRhRu, determining their stability and the short-range order behavior of their constituents.

DOI: [10.1103/PhysRevMaterials.7.045802](https://doi.org/10.1103/PhysRevMaterials.7.045802)

### I. INTRODUCTION

Almost 20 years have passed since independent work from the groups of Yeh [1] and Cantor [2] showed that mixing up to 20 metallic elements in roughly equal parts leads to a smaller-than-expected number of distinct phases, with some corresponding to disordered solid solutions of four to six elements. These so-called high-entropy alloys (HEAs) have since become the subject of intense study [3]. On a fundamental level, the observation of the existence of an extended single-phase stability region for alloys with multiple principal components was surprising, and from a technological standpoint it opened up the possibility of designing materials that defy the limitations of conventional metallurgy and alloy engineering [4,5].

Besides their metallurgical and mechanical applications, HEAs have been found to be promising catalysts [6,7], especially in electrocatalysis [8–10]. They can efficiently reduce overpotentials and boost activities for, e.g., water splitting [11–21], the oxygen reduction reaction [15,18,22–24], or the methanol oxidation reaction [22,25–28] while exhibiting very good stability under reaction conditions. These unusual properties are linked to their multielemental character, which gives rise to four core effects [29,30]: the entropy, sluggish diffusion, lattice distortion, and cocktail effect. While the former two enhance the stability, the latter two can explain the high activity in catalysis. First, lattice distortions occur due to atoms being surrounded by atoms of many different atomic radii, leading to stress and strain. This alters the electronic structure of the alloy. For example, the water splitting activity

of a family of AlNiCoIrX ( $X = \text{Mo, Cr, Cu, Nb, V}$ ) is superior to IrO<sub>2</sub> because the lattice distortion leads to shorter Ir–O bonds [14]. Second, the cocktail effect describes unexpected, synergistic effects of the chosen composition. For instance, the non-noble metal HEA CoCrFeMoNi shows activity for the oxygen reduction reaction similar to that of Pt.

From the computational perspective, modeling HEAs poses a number of distinct challenges. The presence of multiple components requires relatively large simulation cells to unveil microstructures or order-disorder behavior, while the sluggish diffusion requires long timescales and accelerated sampling techniques to overcome free-energy barriers to atom diffusion. Chemical complexity makes empirical force fields inaccurate, and sampling issues make explicit electronic-structure calculations prohibitively demanding. As a consequence, the study of HEAs usually relies on on-site cluster expansions [31,32], together with analytical models that allow us to capture the qualitative thermodynamic behavior [33], even though entropic effects beyond configurational ones are known to play an important role [34]. More recently, force fields based on machine learning (ML) have emerged as an alternative approach, allowing us to match the accuracy of first-principles calculations while describing off-lattice distortions and thermal fluctuations [35–38]. However, the majority of ML frameworks for materials modeling exhibit a poor scaling of memory, computation, and data requirements with the number of chemical species, and so simulations this far have been restricted to a specific combination of four to five elements. In this paper, we introduce a general-purpose ML model for the study of bulk HEAs that uses a recently

proposed strategy to reduce the dimensionality of chemical space, allowing us to generate an accurate and transferable ML potential that can describe arbitrary mixtures of 25 transition metals. The functional form of the model lends itself to an intuitive interpretation of the relations between different transition metals, and careful validation shows that it is capable of accuracy comparable to that of electronic-structure methods in several reference calculations despite the breadth of chemical space it covers. We use this potential to reproduce computationally the seminal Cantor experiments on the decomposition of multielement mixtures, and find a qualitative behavior in the affinity between different species that is consistent with well-known HEAs, allowing us to introduce a data-driven version of the Hume-Rothery rules to guide alloy design. We conclude by studying three alloy compositions: the prototypical Cantor alloy CoCrFeMnNi, its Mn  $\rightarrow$  Mo counterpart that has enhanced catalytic performance, and PdPtIrRuRh—another promising composition for catalysis. In all cases, we observe a tendency to phase-separate at low temperature, and that the short-range order (SRO) observed in high-temperature conditions is indicative of the thermodynamic drive to de-mix.

## II. ALCHEMICAL COMPRESSION OF ML REPRESENTATIONS

We follow the approach introduced in Ref. [39] to reduce the computation, memory, and data requirements of a ML model for a chemically diverse problem. Here we only give a brief overview to highlight the key ideas and introduce the notation. The framework relies on the atom-centered density correlation framework [40], which encompasses most of the widespread descriptors for atomic-scale ML and that is essentially equivalent to the moment tensor potentials [41] and the atomic cluster expansion [42]. The reader is invited to read Ref. [43] (especially Secs. 3 and 7.3) for a more pedagogic discussion. The essential ingredient in this framework is the expansion of the neighbor density within an environment  $A_i$  that describes the atoms in structure  $A$  within a spherical region centered on the  $i$ th atom [Fig. 1(a)], on a basis of radial functions  $R_{nl}$  and spherical harmonics  $Y_l^m$ :

$$\begin{aligned} \langle anlm|\rho_i\rangle &= \int d\mathbf{x} R_{nl}(x) Y_l^m(\hat{\mathbf{x}}) \langle a\mathbf{x}|\rho_i\rangle, \\ \langle a\mathbf{x}|\rho_i\rangle &\equiv \sum_{j \in A_i} \delta_{aj} g(\mathbf{x} - \mathbf{r}_{ji}). \end{aligned} \quad (1)$$

In this expression,  $j$  runs over the neighbors of atom  $i$ ,  $g(\mathbf{x})$  is a Gaussian function (or its Dirac- $\delta$  limit),  $\mathbf{r}_{ji} \equiv \mathbf{r}_j - \mathbf{r}_i$  is the interatomic distance vector between the  $j$ th and  $i$ th atom,  $R_{nl}$  enumerates the radial functions, and  $Y_l^m$  the spherical harmonics. The ket  $|\rho_i\rangle$  indicates the  $i$ -centered neighbor density and the  $a$  index identifies the chemical nature of the atoms.

The bra-ket notation serves to emphasize the fact that the discrete coefficients are simply a projection on a basis of the very same quantity as the real-space neighbor density. We also express the density coefficients with the alternative notation  $\langle an|\rho_i^{\otimes 1}; \lambda\mu\rangle \equiv \langle an\lambda\mu|\rho_i\rangle$ , in which  $\rho_i^{\otimes \nu}$  indicates we are describing the  $\nu$ -neighbor density correlations, and the angular indices ( $\lambda, \mu$ ) are moved to the ket to highlight that they determine the symmetry of the coefficients with

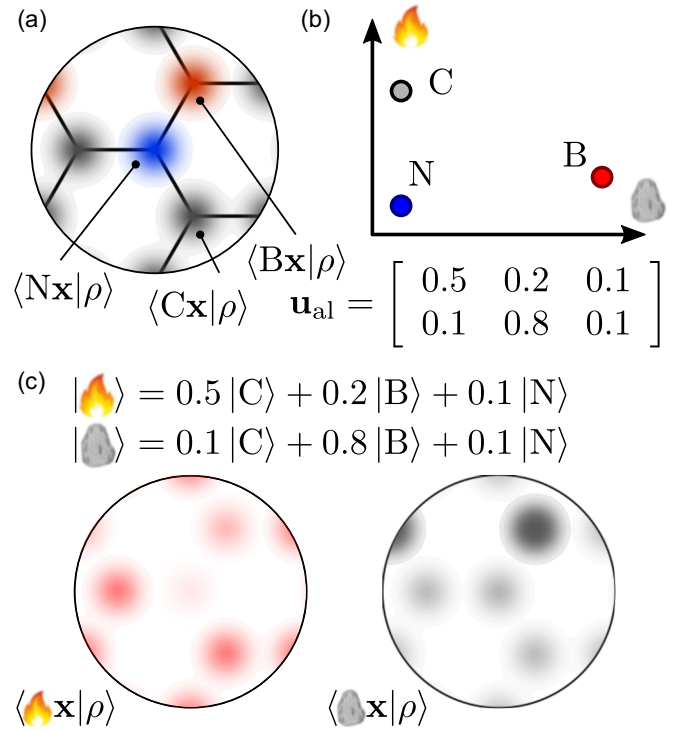


FIG. 1. Different interpretations of the alchemical compression scheme. (a) In a conventional density-correlation ML scheme, each type of atom is associated with a separate density. (b) The entries in the alchemical compression matrix  $\mathbf{u}_{\text{alch}}$  can be interpreted as describing the character of each physical element in terms of  $n_{\text{alch}}$  pseudoelements—a concept that is not dissimilar from the notion of classical elements. (c) The structure can also be seen as described in terms of a density of pseudoelements for which each site contains a contribution from each of the compressed channels.

respect to rotations, that is crucial when building equivariant models and when combining density coefficients to evaluate higher-order correlations.

In this paper, we will use the pair invariants ( $\nu = 1$ ,  $\lambda = 0$ , i.e.,  $\langle an|\rho_i^{\otimes 1}\rangle \equiv \langle an|\rho_i^{\otimes 1}; 00\rangle$ ), as well as the two-neighbor invariant terms ( $\nu = 2$ ,  $\lambda = 0$ ). The two-neighbor invariants—equivalent to Smooth Overlap of Atomic Positions (SOAP) features [44] and closely related to three-body Behler-Parrinello symmetry functions [45]—can be computed as

$$\langle a_1 n_1; a_2 n_2; l|\rho_i^{\otimes 2}\rangle \propto \sum_m \langle a_1 n_1|\rho_i^{\otimes 1}; lm\rangle \langle a_2 n_2|\rho_i^{\otimes 1}; lm\rangle. \quad (2)$$

For readers familiar with the notation used in the SOAP literature [44], the expansion coefficients of the density are often written as  $c_{nlm}^a$  and the power spectrum, corresponding to the two-point correlations (2), as  $P_{n_1 n_2 l}^{a_1 a_2}$ . In both forms, it is clear that the number of components one has to consider grows quadratically with the number of species, each element being considered independently in the neighbor density. The generalization to higher- $\nu$  correlations leads to an even steeper increase, but for the data set we consider here, the computational cost is prohibitive even for two-neighbor correlations.

The key insight in Ref. [39] is that it is unnecessary—and possibly detrimental—to consider elements as independent. Similarities in the behavior of elements have inspired the construction in the periodic table [46] and are routinely used to inform materials design and optimization. Instead, elements should be mapped to a continuous  $n_{\text{alch}}$ -dimensional space, where each chemical species is mapped to  $n_{\text{alch}}$  *pseudospecies* with a set of coupling coefficients  $\mathbf{u}_{\text{alch}}$ . Then, the density coefficients can be contracted as

$$\langle bn | \tilde{\rho}_i^{\otimes 1}; \lambda \mu \rangle \equiv \sum_a u_{ba} \langle an | \rho_i^{\otimes 1}; \lambda \mu \rangle, \quad (3)$$

where we use  $\tilde{\rho}$  to indicate the alchemically compressed neighbor density (Fig. 1). We note that similar ideas were applied—without optimizing the contraction coefficients—in the context of atom-centered symmetry functions [47,48], and that a systematic, rather than data-driven, compression has also been recently applied to an eight-element alloy system in the context of atomic cluster expansion potentials [49]. Moreover, there is a large design space of variations on a theme: separate coupling coefficients could be used depending on the angular ( $\lambda$ ) and/or radial ( $n$ ) channel, and it would be possible to jointly contract over chemical and radial components—which were shown to be effective in reducing the number of features with minimal information loss [50]. Here we do not explore this design space because, as we shall see, the pure alchemical contraction appears to be both effective and easy to interpret. Using these compressed density coefficients (3), one can evaluate correlation functions with a cost that still scales exponentially with  $\nu$ , but with a more benign base, or perform further iterative contraction steps as in Ref. [51].

To conclude this overview, we note that the alchemical coefficients  $\mathbf{u}_{\text{alch}}$  enter the expression for the  $\nu = 2$  features in a quadratic fashion, and so they cannot be directly determined using linear algebra, even if one uses a linear model based on the contracted features. In Ref. [39], this issue was tackled with an iterative strategy, alternating a solution of the linear problem with fixed  $\mathbf{u}_{\text{alch}}$  and a gradient descent on the coupling coefficients. In the present paper, instead, we implement the model using the PYTORCH framework [52], allowing us to use automatic differentiation and gradient descent to optimize simultaneously  $\mathbf{u}_{\text{alch}}$  and the model weights.

### III. COMPUTATIONAL DETAILS

We provide a concise summary of the details of the calculations we perform in this paper, covering the reference electronic-structure calculations, the construction of the training set, and the architecture of the ML model, as well as the details of the sampling protocol that we use for simulations in Secs. VI and VII. In the Supplemental Material [53], we provide representative examples of the typical simulation setup and additional convergence tests.

#### A. Electronic-structure details

All the reference energies and forces are computed using density-functional theory (DFT), as implemented in the VASP code [54], with the PBEsol exchange-correlation functional [55]. The core electrons are treated implicitly using

projector augmented wave (PAW) pseudopotentials [56]. We choose conservative values for the convergence parameters of the electronic structure calculation (see the Supplemental Material [53] for details): the wave function is expanded in plane waves with a cutoff energy of 550 eV, and the Brillouin zone sampling uses a  $\Gamma$ -centered Monkhorst-Pack scheme [57] with an interval between  $k$  points along reciprocal lattice vector  $0.04 \pi \text{ \AA}^{-1}$ . Even though transition metals often exhibit magnetism, either in the pure phases or in alloys, we perform all our calculations without spin polarization. Even disregarding the fact that ML models that can deal with magnetism are still at a very early stage [58], one should consider that we aim to cover a broad chemical range that includes materials which require different types of approaches to describe accurately their magnetic behavior—band magnetism within the local spin density approximation [59], noncollinear magnetism [60], Hubbard- $U$  calculations [61], etc. This makes nonpolarized calculations a reasonable approximation within the scope of the present paper (see also the Supplemental Material [53]), even though this limits the accuracy of our reference and our model for magnetic systems—which, for example, would not be able to predict the stabilization of *bcc* iron over the close-packed polymorphs.

#### B. Training set construction

We generated an original data set including 25 *d*-block elements, i.e., all transition metals excluding those that are not listed in Ref. [62] as relevant for HEAs (Tc, Cd, Re, Os, Hg). We generate a total of 25 000 structures, following a protocol that ensures quasirandom sampling of this high-dimensional phase space. We created four subsets of structures based on *bcc* and *fcc* lattices containing 36 or 48 atoms, respectively. All lattice parameters are defined by the average atomic volume of the elements in a structure and scaled up or down by up to 10% at random to simulate compression and expansion. The structures in the first three classes include from 3 to 8 randomly selected elements, and in the fourth from 3 to 25. In the first class, we included only perfect crystal structures with random compositions. For the three remaining classes, we shuffled atomic positions around their ideal lattice sites (using a Gaussian distribution of atomic displacement with a standard deviation of 0.2  $\text{\AA}$  in the second and fourth classes and 0.5  $\text{\AA}$  in the third) to incorporate the information about interactions in crystals at finite temperatures.

For every class of structures, we generated 100000 random configurations and selected around 7000 of the most diverse from every subset using farthest point sampling [63] in radial spectrum feature space.

#### C. Machine-learning model

We build ML models based on density-correlation representations, combining an atomic-energy baseline, ridge regression based on pair and three-body correlation features, and a multilayer perceptron [64] based on the three-body features. Here we discuss briefly the functional form of the different term and outline the training strategy we followed. The atomic-energy baseline is simply a linear model that depends exclusively on the nature of the atom at the center

of each environment,  $a_i$ :

$$V^{(\text{aeb})}(A_i) = w_{a_i}^{(\text{aeb})}. \quad (4)$$

Even though we train on atomization energies (and so the large dependency of the atomic energies on the details of the pseudopotentials is not an issue) we still find that  $V^{(\text{aeb})}$  captures a large fraction of the target variance and facilitates learning. The second term we consider is a set of pair energies. We use a Gaussian width of 0.25 Å, a cutoff of 6 Å, and radial scaling following Ref. [39]; we expand the density in spherical harmonics and in 12 radial functions, enumerated by the  $n$  index and obtained by orthogonalizing Gaussian-type orbitals that cover the range of distances up to the cutoff radius (see, e.g., Ref. [65] for a precise definition). We use different weights depending on the nature of the two atoms, so in practice the contribution to the potential reads

$$V^{(2\text{B})}(A_i) = \sum_{an} w_{ai,an}^{(2\text{B})} \langle an | \bar{\rho}_i^{\otimes 1} \rangle. \quad (5)$$

The third term involves three-body correlations (SOAP features), computed on top of algebraically contracted density coefficients, with a linear model

$$V^{(3\text{B})}(A_i) = \sum_{bnb'n'l} w_{bnb'n'l}^{(3\text{B})} \langle bnb'n'l | \bar{\rho}_i^{\otimes 2} \rangle. \quad (6)$$

We use the same set of weights irrespective of the atom type, because in a three-body descriptor the nature of the central atom is encoded in the density associated with the Gaussian at  $r = 0$ , so the compression of the dependency of potentials on the central atom type is achieved implicitly and with the same contraction coefficients used for the neighbor density.

Finally, we include a nonlinear term that takes the compressed power spectrum as input and feeds it into a Behler-Parrinello-style [66] multilayer perceptron [64]. First, a linear filter projects the power-spectrum features into 80 input neurons,  $\xi^{(0)}$ , to which hyperbolic tangent activation functions are applied. A second linear layer combines the outputs of the neurons, feeding them to one hidden layer of the same size. Finally, the outputs are linearly combined to yield the atomic energy:

$$\xi_q^{(0)}(A_i) = \sum_{bnb'n'l} w_{qbnb'n'l}^{(\text{NN},0)} \langle bnb'n'l | \bar{\rho}_i^{\otimes 2} \rangle, \quad (7)$$

$$V^{(\text{NN})}(A_i) = F(\xi^{(0)}(A_i)).$$

We use this simple Neural Network (NN)—built on top of the compressed power-spectrum features—because we want a simple and well-understood term that can incorporate non-linearity without exploding the design space, and because we want to show that our alchemical compression scheme can be readily applied to several well-established ML schemes. It is possible (and likely) that alternative frameworks e.g., increasing further the body order, may allow for a better-performing model, but as we shall see this approach is sufficient to achieve state-of-the-art accuracy together with a stable and interpretable model.

The parameters of  $V^{(3\text{B})}$  and  $V^{(\text{NN})}$  implicitly include the alchemical coupling matrix  $\mathbf{u}_{\text{alch}}$ ; for this reason, we optimize

all models with gradient descent, relying on backpropagation as implemented in PYTORCH [52]. A ridge penalty term is included on all weights to reduce the risk of overfitting. We find that (possibly due to the presence of large linear components that contribute a quadratic term to the  $L^2$  loss) a deterministic L-BFGS optimizer [67] performs much better than stochastic gradient descent.

#### D. Sampling details

Molecular dynamics (MD) is well-suited to describe structural relaxation of the atomic coordinates. However, long-range diffusion in the solid phase occurs through vacancies and is too slow to be simulated explicitly by MD. To overcome this timescale problem, we use a combination of techniques to facilitate thorough sampling of atomic ordering. Our base protocol involves performing MD simulations in the constant-temperature/constant-pressure  $NpT$  ensemble [68]. We use a conservative time step of 2 fs, an isotropic barostat [69] with a time constant of 200 fs coupled to an optimal-sampling colored-noise thermostat [70] and an aggressive thermostat for the ions, alternating an optimal-sampling Langevin equation with a stochastic velocity rescaling [71] with a time constant of 10 fs. We accelerate sampling of the compositional (dis)order by performing Monte Carlo (MC) steps in which the nature of two atoms in the system is exchanged, with a Metropolis acceptance criterion [72]. We perform on average one exchange attempt per MD time step. Both the MD and the MC step conserve the Boltzmann distribution (except for a negligible finite time-step error), and so the combined MD/MC protocol is consistent with canonical sampling. To further accelerate sampling, we also use replica exchange molecular dynamics (REMD) [73]—a technique in which multiple trajectories at different temperatures are performed in parallel. Periodically, structures are exchanged between temperatures, using a Monte Carlo procedure that preserves the Boltzmann distribution for each thermodynamic state. The fact that each trajectory is brought through cycles of heating and annealing accelerates conformational sampling and reduces the correlation time of observables that are associated with activated events at low temperature. Unless otherwise specified, we use temperature replicas distributed according to a geometric progression between two extremal values  $T_{\min}$  and  $T_{\max}$ . For all MD/MC simulations, we use the i-PI universal force engine [74] that includes an implementation of element exchange moves [75] and a flexible implementation of replica exchange [76].

## IV. ALCHEMICAL LEARNING

As discussed above, the compression scheme in Eq. (3) is just one of the many approaches one could take to reduce the dimensionality of the density expansion coefficients. One of the appealing features of this specific implementation is that it can be interpreted relatively easily, and that it allows us to extract physical-chemical insights through an introspection of the model parameters and performance.






FIG. 2. Learning curves for different models. Full lines correspond to models built using only  $V^{(aeb)}$  and  $V^{(3B)}$ , with  $n_{alch}$  pseudoelements (all optimized iteratively). The dotted green curves are obtained with a  $\mathbf{u}_{alch}$  filled with uniform random numbers (rnd.) and with the weights we use as an initial guess for the optimized models (base) that are built based on physical priors following the scheme discussed in Ref. [39]. The dashed green line corresponds to a model that includes  $V^{(aeb)}$  and  $V^{(3B)}$ , as well as the full set of pair potentials and a nonlinear term built on top of the contracted power spectrum features  $V^{(NN)}$ .

### A. Learning curve analysis

We begin by considering linear models based on contracted power-spectrum features, supplemented by an atomic energy baseline term,  $V^{(aeb)} + V^{(3B)}$ . We perform separate training exercises, using only energy as targets, and restricting the alchemical contraction to two to five pseudoelements. For each model, we compute learning curves by converging the loss at a given number of training structures  $n_{train}$ , then increase the train size and continue the optimization restarting from the previous weights. Given that the optimization procedure is rather demanding, we do not perform multiple train/test splits, but use consistently the same shuffle with up to 25 000 structures used for training and a holdout set containing 500 configurations used for testing. Even though the accuracy does depend slightly on the shuffle, and on the initialization of the weights, we find that the qualitative observations we present here are robust.

Figure 2 shows a behavior similar to that observed in Ref. [39] for an analogous exercise on the elpasolites data set [77]: at the smaller train set sizes, a very aggressive compression is effective at obtaining a robust model, but with more training data the learning curves saturate. Increasing the number of pseudoelements  $n_{alch}$  delays saturation, but the improvement going from  $n_{alch} = 3$  to  $n_{alch} = 4$  is negligible, and the learning curves for  $n_{alch} = 5$  sit almost exactly at the same value. This indicates that, from the point of view of three-body interactions, three to four pseudoelements are sufficient to saturate the descriptive power of a linear model. Note that the optimization of  $\mathbf{u}_{alch}$  is critical to achieve such efficient compression: a model that uses fixed, random values for the contraction weights, as well as one that uses a fixed, physically inspired initialization of  $\mathbf{u}_{alch}$ , lead to an order of

magnitude increase in the saturation error, even with  $n_{alch} = 4$  (Fig. 2).

Given the saturation of  $V^{(3B)}$ , we proceed to increase the effective body order of the potential adding a nonlinear NN layer on top of the contracted power spectrum,  $V^{(NN)}$ , which introduces about 160 000 additional model parameters, mostly associated with the contraction of the  $|\tilde{\rho}_i^{\otimes 2}\rangle$  features to the 80 input features of the NN. Furthermore, we also include a noncompressed two-body potential  $V^{(2B)}$ , for which we also consider a slightly larger cutoff distance. This two-body term, on its own, does not improve significantly the limiting accuracy of the model (reinforcing the notion that the alchemical contraction is converged) but we include it because it is inexpensive to compute, and has been shown in the past to lead to more stable models, whose performance degrades more gently in the extrapolative regime [78,79]. Incorporating a nonlinear term in the model allows us to overcome the saturation of the learning curve (Fig. 2, dashed green line). The nonlinear  $n_{alch} = 4$  model reaches a validation-set mean absolute error (MAE) below 10 meV/atom. We discuss further the accuracy of this model (that we will refer to as the HEA25-4-NN) in Sec. V.

### B. A 3D periodic table for the transition metals

The alchemical coupling matrix associates to each of the physical elements a vector of size  $n_{alch}$  that can be regarded as the composition of that element in terms of a set of pseudoelements [Fig. 1(b)]. Thus, different atomic species can be seen as points in a continuum space and can be visualized as such to gain insights into the data-driven similarities that arise from the optimization of  $\mathbf{u}_{alch}$  to achieve the most accurate regression of the target. To make the visualization independent on unitary transformations of the weight matrix, we perform a principal component analysis.

The eigenvalues of the covariance matrix indicate the magnitude of the various components (their *explained variance*), and provide another indication of the importance of successive increases in the dimensionality of the alchemical space. We observe a quick decrease of the explained variance, with the fourth component typically amounting to less than 2% of the variance (Fig. 3, inset). This confirms that the first three components provide sufficient descriptive power to capture the difference in behavior between transition metals. We can then look at how the  $d$ -block elements appear when projected along the top three principal components of  $\mathbf{u}_{alch}$  (Fig. 3). We focus on the weights from the HEA25-4-NN model, but the qualitative features of the alchemical projections are similar also for other models in Fig. 2 (see the Supplemental Material [53]). The elements are arranged in a way that is strongly reminiscent of their placement in the  $d$  block: the third principal direction corresponds to the period, while the first two dimensions are associated with a semicircular arrangement, with the elements appearing in the same order as the columns in the conventional periodic table. Interestingly, this arrangement is reminiscent of that used for the  $d$  block in some of the alternative representations of the periodic table, such as the Benfey spiral [80]. It indicates that, from the point of view of the construction of an interatomic potential, zinc is closer to

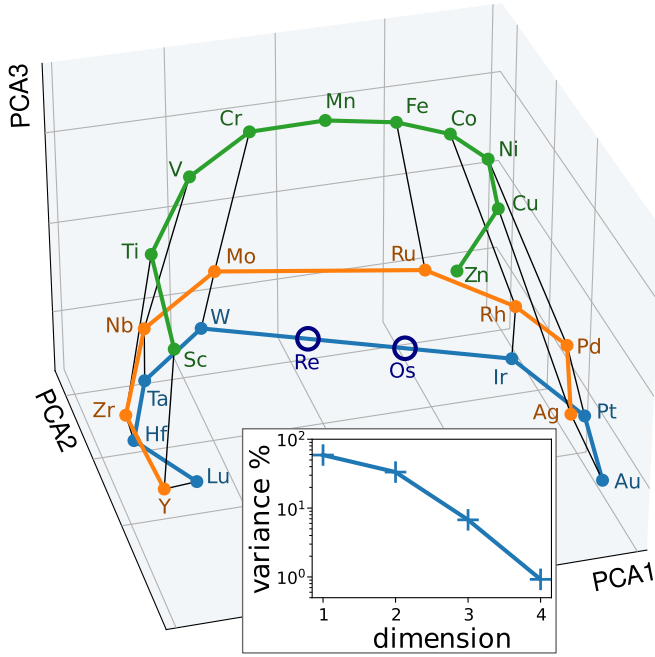


FIG. 3. Top-three principal components of the alchemical coupling matrix  $\mathbf{u}_{\text{alch}}$  for the HEA25-4-NN model. The periods are highlighted with orange, blue, and green lines, and the columns are indicated by black thin lines. Interpolated positions for Re and Os are indicated with empty circles. The inset shows the decay of the explained variance for the four principal components.

scandium then it is to the atoms in the middle of the transition metals block.

### C. Alchemical interpolation

The elements we have not considered leave a clear gap in the arrangement of the alchemical coupling weights, and it is interesting to see how accurate a model that places rhenium and osmium between tungsten and iridium fares in predicting their properties without additional fitting.

We pick 60 structures from the holdout set, containing distorted configurations with random composition. The MAE for these structures when using the  $n_{\text{alch}} = 4$  model using only  $V^{(\text{aeb})}$  and  $V^{(3B)}$  is 13 meV/atom. We then substitute some random atoms with Re and Os, without changing the positions, and recompute their energies with analogous DFT settings.

We then build a model in which we simply take the parameters optimized for the 25-element data set and complete them by adding atomic-energy baselines for Re and Os (obtained by training on the residual a two-parameter model that depends exclusively on the Re and Os content) and by adding pseudo-element weights that interpolate linearly between W and Ir (see Fig. 3):

$$u_{b\text{Re}} = \frac{2}{3}u_{b\text{W}} + \frac{1}{3}u_{b\text{Ir}}, \quad u_{b\text{Os}} = \frac{1}{3}u_{b\text{W}} + \frac{2}{3}u_{b\text{Ir}}. \quad (8)$$

The power-spectrum model weights are unchanged: we are effectively interpolating in pseudo-element space. The resulting model yields exactly the same predictions for structures that do not contain Os and Re, and has a MAE of only 24

meV/atom for the test structures that include the two species (see also the Supplemental Material [53]). The model is also sufficiently stable to run MD simulations for Re and Os containing structures.

This example underscores the advantages of the interpretable functional form we use to implement alchemical dimensionality reduction. It also opens up the possibility of designing simulation protocols that include smooth alchemical transformations in a similar spirit as the framework pioneered by Sheppard *et al.* [81], and with some similarities to the virtual crystal approximation that is often used to describe approximately random alloys [82]. For example, one could use thermodynamic integration to compute the change in chemical potential associated with an element substitution by running simulations with a mixed potential, in which the alchemical coupling weights are gradually transformed between the values associated with two elements.

## V. VALIDATION OF THE POTENTIAL

We now assess the accuracy and stability of the model we use in the rest of this paper, which combines a four-pseudoelement contraction of the power spectrum with a multilayer perceptron. We aim to provide benchmarks that are easy to reproduce but that reflect the performance of the model in relevant simulation tasks, and we envisage that any comparative study would include most of these and not only cross-validation statistics. To contextualize and provide a reference scale for our results, we report in the Supplemental Material [53] similar validation results for the general-purpose, universal graph neural network M3GNet [83]. In all cases, HEA25-4-NN, which admittedly has a narrower scope of applicability, outperforms M3GNet by a large margin.

### A. Holdout validation of the HEA25-4-NN model

We train the HEA25-4-NN potential by progressively increasing the train set size until we run the final optimization on 25 000 structures, including forces for 2000 of them. We hold out 500 structures and use them for validation. The parity plot between targets and predictions demonstrates the accuracy of the model (Fig. 4), which is remarkable given the diversity of the data set, which contains random combinations of up to 25 elements, and highly distorted structures.

### B. Binary convex hulls

Even though the HEA25-4-NN is clearly geared towards multicomponent simulations, it is important that it also provides reasonable results for simpler compositions, as these may appear spontaneously when complex alloys de-mix and form precipitates. We collect 1438 binary intermetallic structures out of more than 146k crystal structures from the Materials Project database [84] and recompute their energies with single-point calculations using our DFT setup, as well as with the HEA25-4-NN model. We discard 23 structures for which our DFT calculations did not converge and 10 that correspond to configurations that are too dissimilar from the bulk structures we consider here (see Supplemental Material




FIG. 4. Parity plot between reference energy and forces and the values computed with the HEA25-4-NN model for a holdout set of 500 structures, randomly selected from the overall pool of structures. Energy error: 10 meV/atom mean absolute error (MAE), 14 meV/atom root mean square error (RMSE), force error: 190 meV/Å MAE, 280 meV/Å RMSE.

[53]). For the remaining structures, the MAE error for the cohesive energy is 62 meV/at. and for the formation energies is 63 meV/at, which is higher than the cross-validation error, but still remarkably accurate for extrapolative predictions. It is worth noting that the MAE discrepancy between our DFT calculations and those saved in the MP records is 65 meV/at.; this is due to the significant difference in the details of the electronic structure calculations, e.g., the use of Hubbard U corrections for some structures in the MP protocol, and neglect of spin polarization in ours. This observation underscores that the details of the electronic structure calculations can have an impact comparable to the accuracy of our ML model. We then use this data to compute binary convex-hull diagrams for all element pairs. In Fig. 5, we show a representative example for the Ti–Pt system. The overall shape of the hull is usually well-reproduced, but often HEA25-4-NN predicts different stable polymorphs than DFT and/or mispredicts the stability of certain compositions (as is the case for  $\text{TiPt}_2$  in the figure). However, these qualitative errors are usually associated with situations in which a small energy shift can bring a composition above the hull boundary, and even in a fully *ab initio* study it would not be possible to determine conclusively its thermodynamic stability. The full list of hulls is included in the Supplemental Material [53]. Figure 5 also shows an overview of the accuracy of the prediction of formation energies for all phases (stable and unstable) as a function of composition. Errors are not uniform: some elements such as Mn, that have the tendency of forming complex crystal structures, yield larger errors, while others such as Cu or Ni usually yield errors comparable to the validation set. It would be trivial to improve the accuracy of the model for binary structures and pure element polymorphs by including


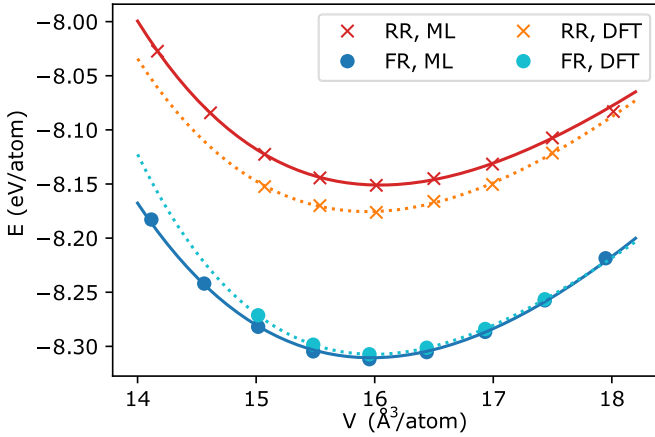


FIG. 5. MAE for the formation energy of binary compounds from the Materials Project database. The inset shows a representative hull plot for the Ti–Pt system, highlighting the hulls obtained from the single-point DFT calculations and the ML predictions. The dashed line identifies the structures that are stable based on the energies available in the Materials Project database.

this small number of additional structures in the training set. We chose not to do that to avoid introducing biases in the accuracy depending on the different abundance of structures in the MP database. In the future, we plan to systematically extend our training set to incorporate disordered and liquid structures.

### C. Energy and equation of state

We prepare a  $5 \times 5 \times 5$  *fcc* supercell, containing five atoms of each of the 25 elements, arranged randomly on the lattice. We relax the geometry of the structure and the volume of the supercell using the HEA25-4-NN potential. We refer to this structure as the random relaxed (RR) structure. Starting from the same configuration, we also perform a slow annealing trajectory, combining MD and atom exchange moves, to obtain a structure in which the arrangement of elements is not random but more energetically favorable. We refer to this structure as the fully relaxed (FR) structure. In both cases, the atoms relax away from *fcc* lattice positions, and the resulting structure within the supercell is rather disordered. We then introduce an isotropic compression or expansion of the two structures, relaxing the coordinates of the atoms within the cell, and fit a Birch-Murnaghan equation of state to the resulting energy-volume curves. We repeat the fixed-cell relaxation with the reference DFT and compare the resulting equations of state (Fig. 6). The error on the cohesive energy  $E_0$  is comparable to the test error (24meV for  $E_0^{(RR)}$ , 3meV



	RR, ML	RR, DFT	FR, ML	FR, DFT
$E_0$ (eV/at.)	-8.151	-8.175	-8.310	-8.307
$V_0$ ( $\text{\AA}^3/\text{at.}$ )	16.04	15.94	15.97	16.00
$B_0$ (GPa)	132	141	146	162
$B'_0$	6.7	6.1	4.8	8.0

FIG. 6. Equation of state for the random relaxed (RR) and fully relaxed (FR) structures (see text for the full definition), computed with the HEA25-4-NN potential and with the reference DFT. Birch-Murnaghan parameters for cohesive energy ( $E_0$ ), equilibrium volume ( $V_0$ ), bulk modulus ( $B_0$ ), and bulk modulus derivative ( $B'_0$ ) are given in the table.

for  $E_0^{(\text{FR})}$ ) and much smaller than the energy gain associated with the annealing of the lattice occupations ( $E_0^{(\text{RR})} - E_0^{(\text{FR})}$  is about 150 meV/atom), indicating that HEA25-4-NN is reliable for assessing the energetics of ordering in a random alloy. The equilibrium volume and bulk modulus for the two structures are also in good agreement, with errors below 1% and 10%, respectively—comparable to the typical discrepancy between different DFT approximations or between DFT and experiments.

#### D. Molecular dynamics

As a further demonstration of the accuracy and the stability of this potential, we perform two constant-pressure MD/MC trajectories, one at  $T = 300$  K and one at  $T = 5000$  K, each starting from a random arrangement of five atoms for each of the 25 elements (a total of 125 atoms) arranged on an *fcc* lattice. The trajectories are 10 ps long, with on average one attempt at exchanging a pair of atoms every 2 fs. We save a configuration every 100 fs, and perform DFT calculations to compare energy and forces with those obtained from the ML potential. Figure 7 shows that the low-temperature trajectory, where major rearrangements of the atoms occur but the structure remains approximately *fcc*, has an accuracy comparable to that measured on the validation set. The high-temperature run exhibits a higher error. However, the main component of the error is a rigid shift of the energies, and the trajectory remains stable—which is remarkable given that we observe complete melting, and the potential is trained exclusively on distorted solid structures.

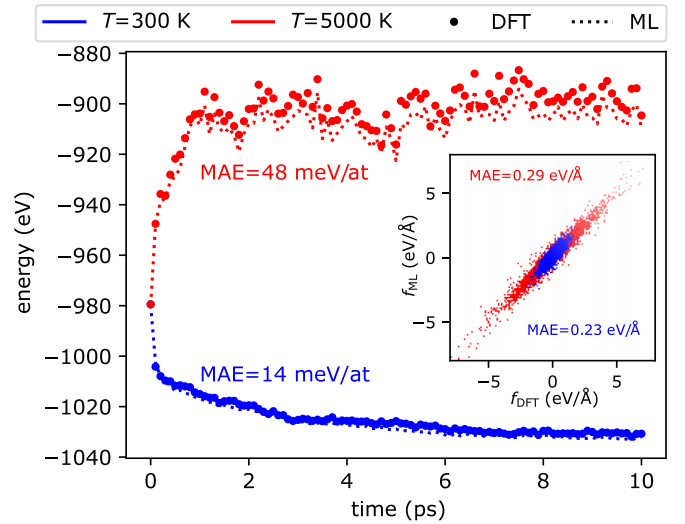


FIG. 7. Comparison between the potential energy evaluated along two 10 ps MD/MC trajectories and that recomputed by DFT for 100 snapshots. The inset shows the parity plot for the force components computed for those structures. Energies have a MAE of 14 (48) meV/atom and forces a component MAE of 0.23 (0.29) eV/ $\text{\AA}$  for the 300 (5000) K trajectory.

## VI. TEMPERATURE-DEPENDENT SEGREGATION IN A CANTOR-STYLE ALLOY

In a seminal experiment, Cantor *et al.* [2] investigated the development of microstructures during the solidification of equimolar mixtures of 16 and 20 elements. We aim to perform a similar experiment in a computational setting, assessing the propensity of different elements to pair together or segregate, while covering the full component palette allowed by our model. This poses considerable challenges beyond the chemical complexity: kinetic trapping plays an important role in the physics of HEAs, and simulating vacancy-assisted atom diffusion requires timescales that are unattainable in brute-force atomistic modeling. To accelerate sampling and achieve (partial) equilibration, we run replica exchange simulations combining MD and atom swap moves (REMD/MC), as described in Sec. III D.

Figure 8 shows a representative trajectory for an 864-atom cell, starting from *fcc* configurations, and including equimolar composition of all 25 elements (a composition we will refer to as HEA<sub>all</sub>). The slow, logarithmic relaxation of the low-temperature replica is indicative of the glassy dynamics of the system, which does not equilibrate completely even after millions of MD/MC steps (see the Supplemental Material [53]). For this reason, we perform multiple independent (and longer) simulations with a smaller box size (see the Supplemental Material [53]). The qualitative observations on the local ordering are robust, even though the precise arrangement of atoms in the low-temperature regime, as measured by the element-resolved pair correlation functions, differ noticeably between trajectories.

#### A. Relative pair probabilities for the HEA<sub>all</sub> alloy

The pair correlation functions (Fig. 9) display broad, liquidlike peaks at both the highest and lowest temperatures we






FIG. 8. Trajectories of the potential energy for the 40 replicas used in one of the REMD simulations of an 864-atom box of the HEA<sub>all</sub>. Each color corresponds to a different initial configuration, that goes through cycles of heating and cooling due to REMD exchanges, accelerating the equilibration of the simulation at each temperature. The collection of trajectory segments corresponding to the extremal temperatures  $T = 300$  K and  $T = 1253$  K are highlighted with thicker, black lines. The logarithmic timescale refers to the MD integration time, but should not be interpreted as physical time given the presence of MC steps and replica exchange moves.

considered. In fact, simulations show little diffusion (except for some occasional bursts of activity at the high end of the temperature range) and the system can be characterized as an amorphous (or nanocrystalline) solid. The broadening of the peaks can be at least in part attributed to the diversity of pair distances between atomic species: some, like Cr-Cr, peak at distances as short as 2 Å, others, such as Y-Y, peak at about 3.7 Å. Note that typical distances in same-element pairs do not always match those found in the pure solid, underscoring the fact that the HEA25-4-NN can capture the effects arising from the heterogeneous chemical environments found in this alloy. For this reason, and given the disordered structure that develops in the supercell, we analyze structural correlations using a coarse-grained definition in which the first coordination shell extends up to a distance  $r = 3.75$  Å, the second up to  $r = 6.25$  Å, and the third up to  $r = 8$  Å, which is the largest




FIG. 9. Pair correlation functions computed on a the  $T = 300$  K (full) and  $T = 1253$  K (dashed lines) replicas of a HEA<sub>all</sub> box. Black lines correspond to the unresolved pair correlation, while red (Cr-Cr) and blue (Y-Y) lines provide representative examples of pair correlations resolved by species. The vertical dotted lines indicate the regions used in the definition of the pair ordering.

distance we consider given the size of the box. We then define a variation on a theme of the SRO parameter [85], which we dub the relative pair probability (RPP),

$$\text{RPP}_{\Delta r}(A, B) = \frac{p_{\Delta r}(A, B)}{p_{\Delta r}(\star, \star)} \frac{\rho^2}{\rho_A \rho_B}, \quad (9)$$

which computes the number of pairs between species A and B that occur within a range  $\Delta r$  of distances, divided by the number of all pairs found in that same region, and normalized by the number density of the two species,  $\rho_{A,B}$  and the overall number density  $\rho$ .  $\text{RPP} = 1$  indicates that the two species are as likely to be found within a given separation range than any atom pair.  $\text{RPP} > 1$  ( $< 1$ ) indicate that they are more (less) likely to be found in that distance range.

Qualitatively, the value of the RPP in the first coordination shell is indicative of the propensity of two elements to cluster together or to separate from each other. However, the values cannot be interpreted in isolation without considering the overall setup of the simulation: the finite size of the supercell, the imperfect equilibration, and the many-body interactions between all 25 species mean that the strong affinity between Y and Au, or the poor compatibility of Mn and Pd, do not necessarily imply the same quantitative effect when considered as part of a different overall composition. Figure 10 shows a heat-map representation of  $\text{RPP}_{\Delta r}(A, B)$  for the HEA<sub>all</sub> at 300 K and 1253 K, and for the three regions indicated in Fig. 9. A few qualitative observations can be made. First, in our simulations HEA<sub>all</sub> evolves to be far from random. Certain atom pairs have a strong tendency to associate or separate at low temperature, and the high-temperature samples (which are well equilibrated) show similar, even though less pronounced, trends. This correspondence is interesting, as it suggests one may use high-temperature trajectories, which are easier to converge, to extract insights on the propensity of different species for association. The trends observed in the second and third regions are very similar to those in the first-extended-neighbor shell, although progressively less pronounced: given the finite size of the simulation, and incomplete equilibration, the simulation does not generate clear-cut phase-separated regions.

Considering the RPP along the elements, one can observe a clear periodicity in behavior. Sc, Y, Hf, as well as the noble metals, Cu and Zn, tend to separate from V, Cr, Mn, Fe, which on the other hand have a tendency to cluster together, and also have positive associations to their heavier counterparts Nb, Mo, Ta, W. On the other hand, Sc, Y, and (to a lesser degree) Hf associate strongly with noble metals, Cu, and Zn. The noble metals, Cu, and Zn also tend to cluster together. Ti, Co, Ni, Zr, Ru, Ir have less clear-cut associations and are closer to having a random distribution throughout the box. Another way of looking at the association plots in Fig. 10 is to check for consistency with known HEAs. The Cr-Mn-Fe-Co-Ni system is one of the prototypical sets of HEA formers, and indeed we observe strong mutual association tendency between Cr-Mn-Fe in the first shell, and also with Co and Ni in the second extended shell. Second-shell mutual association is also observed for noble-metal-based compositions such as Ni-Cu-Pd-Pt-Au. Let us reiterate that

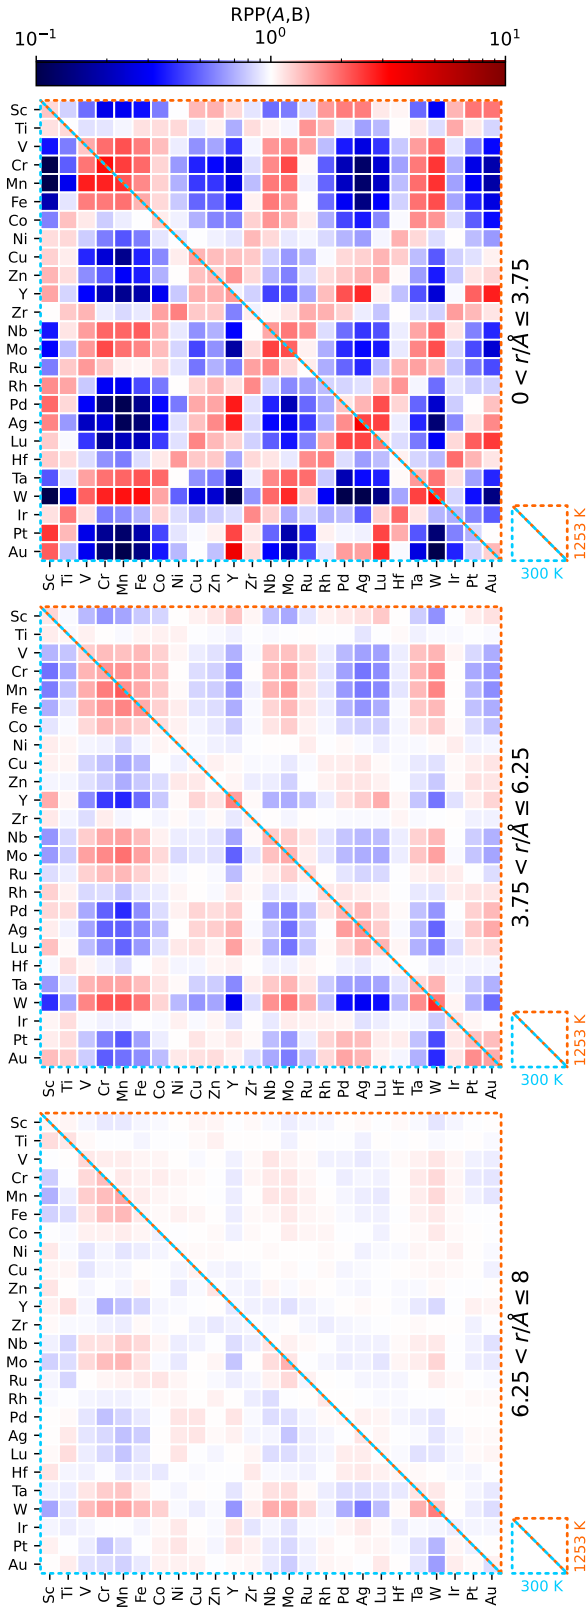


FIG. 10. A plot of the relative pair probability for all atom pairs and the three regions corresponding to the first, second, and third peaks in the total pair correlation function (Fig. 9). Each plot is split across the diagonal, showing results for simulations of HEA<sub>all</sub> at two temperatures, 300 K (lower-left corner) and 1253 K (top-right corner), averaged over the trajectories and discarding the first 100 ps (50 000 combined MD/MC steps).

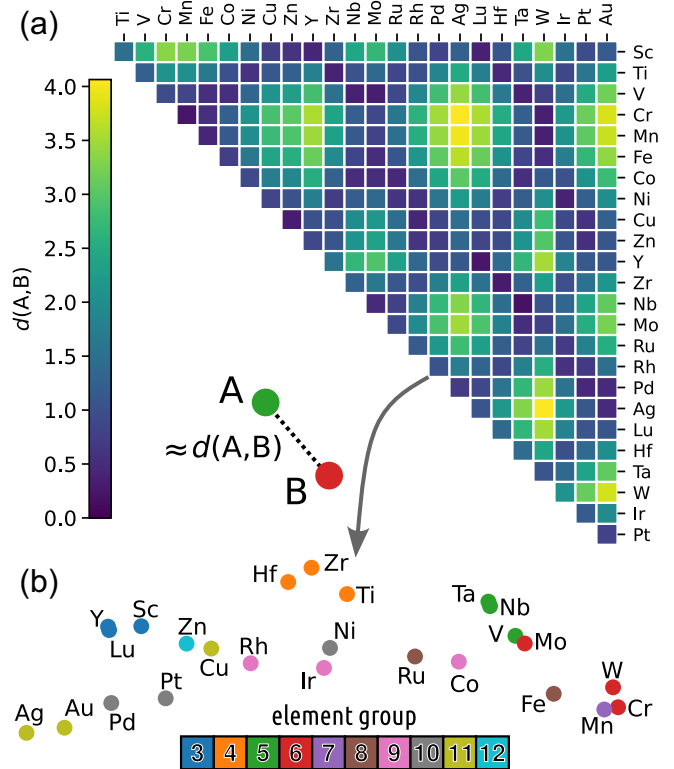


FIG. 11. (a) Element similarity matrix based on the RPP distance (10) for the nearest-neighbor shell, in the HEA<sub>all</sub> simulation at  $T = 1253$  K. (b) The element similarity map (color-coded based on the group of the various transition metals) is built by applying metric multidimensional scaling to the distance matrix and provides a visual aid to recognize groups of elements that have similar affinity patterns to the other  $d$ -block metals.

strong mutual association for a group of elements in the HEA<sub>all</sub> runs is a necessary, but not sufficient, condition for that group of elements to be good HEA-forming candidates. For instance, some elements may have a strong tendency to form ordered intermetallics and might separate out of the mixture.

### B. Data-driven Hume-Rothery rules

This analysis allows us to substantiate and quantify some of the empirical principles that are used in the design of HEAs, such as Hume-Rothery rules [86] that stipulate what elements can be substituted for each other with little effect on the HEA-forming propensity. We use the first-neighbor affinity of each species to all the other elements in the alloy to define a measure of dissimilarity as

$$d_{\text{RPP}}(A,B)^2 = \sum_X \left[ \log_{10} \frac{\text{RPP}_1(A,X)}{\text{RPP}_1(B,X)} \right]^2, \quad (10)$$

that, roughly speaking, measures the relative strength of interactions between the two species and the other components. Two elements with a small distance are predicted to behave similarly, and vice versa. Figure 11 paints a picture that is consistent with the observations we made on short and midrange order between the elements in the HEA<sub>all</sub>, and with much of

the common wisdom in HEA research. We base this analysis on the high-temperature simulations to obtain a statistically-converged, and somewhat more nuanced, definition, but the qualitative features of the map are similar to those one would obtain from the RPP computed at  $T = 300$  K. Elements in the same group usually show strong similarity, but this is not always the case: for example, Cu is more similar to Zn than to Ag. The similarity matrix can also be converted to a 2D map, in which the Euclidean distance between elements approximates their RPP-based similarity (also shown in Fig. 11), which provides an easy-to-interpret visual representation of a set of data-driven rules to design HEAs. The element similarity that can be inferred from the RPP-based map differs—both quantitatively and conceptually—from that associated with the alchemical coupling matrix in Fig. 3. Whereas the weights are associated with the similarity in terms of the interatomic potential, the RPP similarity is a result of the collective behavior of the  $\text{HEA}_{\text{all}}$  at the prescribed thermodynamic conditions, not unlike the relation between a pair potential and the potential of mean force. This means, for example, that one could compute  $d_{\text{RPP}}$  for a different alloy composition (extending or refining the assessment of alloying behavior) from a different type of interatomic potential or even from experimental data on partial structure factors.

## VII. BULK STRUCTURE OF HIGH-ENTROPY ALLOYS FOR CATALYSIS

Having demonstrated the accuracy of the HEA25-4-NN model and having used it to investigate the mutual affinity of the full set of 25 transition metals we considered in a Cantor-type computational experiment, we now turn our attention to a more focused study of three specific equimolar compositions. The first is the prototypical CoCrFeMnNi alloy, which was reported by Cantor *et al.* [2] in their seminal paper. This alloy is also known to be effective as a catalyst [87,88]. Furthermore, we investigate CoCrFeMoNi [11,89,90] as an example of an alloy obtained by element substitution that has been broadly studied for its improved mechanical and tribological properties [91,92], as well as a catalyst of oxygen evolution reactions. We then consider IrPdPtRhRu [21,93–96] as an example of an alloy based on sixth-period elements that has recently received much attention as a catalyst for hydrogen evolution and is often synthesized in the form of nanoparticles.

To model the alloys, we used *fcc* lattices with 500 atoms per cell ( $5 \times 5 \times 5$  supercell). We ran two independent REMD/MC runs according to Sec. III D with a time step of 2 fs and 32 temperature replicas, logarithmically spaced between 300 K and 1253 K. We discard the first 100ps for equilibration. Given that all these alloys maintain a regular *fcc* structure throughout the simulation, we analyze their structure in terms of Cowley's SRO [85], which is commonly used in the study of HEAs and takes a value of zero when atoms are distributed fully randomly, becomes negative for pairs of atoms that tend to cluster together, and tends to one when two atom types never appear as first neighbors. In the Supplemental Material [53], we also report an analysis in terms of the RPP that incorporates second-neighbor and long-range correlations. In interpreting these results, one should keep in mind




FIG. 12. (a) Cowley's short-range (SRO) parameters for the first shell in CoCrFeMnNi HEA, shown for the ten replicas between 300 and 1253 K, averaged over the last 1000 steps and two independent runs. At low temperatures, a tendency of Fe-Mn segregation can be seen. In contrast, Co is very well mixed. There are two phase transformations around 400 K and 900 K. The y axis is adjusted to the example shown in Fig. 14 to facilitate comparison. (b), (c) Snapshot from MC/MD simulations at  $T = 300$  K and at  $T = 720$  K, respectively. In the 300 K snapshot, two planes of Ni can be seen, while in the higher temperature snapshot, Cr order is evident (see the Supplemental Material [53]).

similar considerations to those we discussed for the  $\text{HEA}_{\text{all}}$  simulations: (1) the SRO (and the RPP) are only meaningful for homogeneous phases, and in case of phase separation the values computed for the whole cell serve only to signal the occurrence of a phase transition; (2) a combination of finite-size effects and glassy behavior can hinder reaching full equilibrium in simulations; and (3) since they allow for atom exchanges, our simulations cannot give quantitative indications on whether different phases are only metastable nor on the kinetics of diffusion processes that are only required for precipitation.

We start by analyzing the Cantor alloy CoCrFeMnNi. The SRO computed at different temperatures [Fig. 12(a), plotted for all element combinations] indicate the presence of at least two phase transitions. The high-temperature phase is homogeneous and disordered but shows substantial ordering, particularly for the Cr-Cr pair. At approximately 900 K, we observe a first transition that is associated with the ordering of Cr atoms. The SRO for the Cr-Cr pair tends to one (as there are almost no first-neighbor chromium atoms) but the RPP shows a clear increase of second-neighbor Cr-Cr pairs, consistent with the formation of a simple cubic sublattice. The other elements remain relatively disordered, and no discontinuous behavior is observed in the SRO. As the temperature is



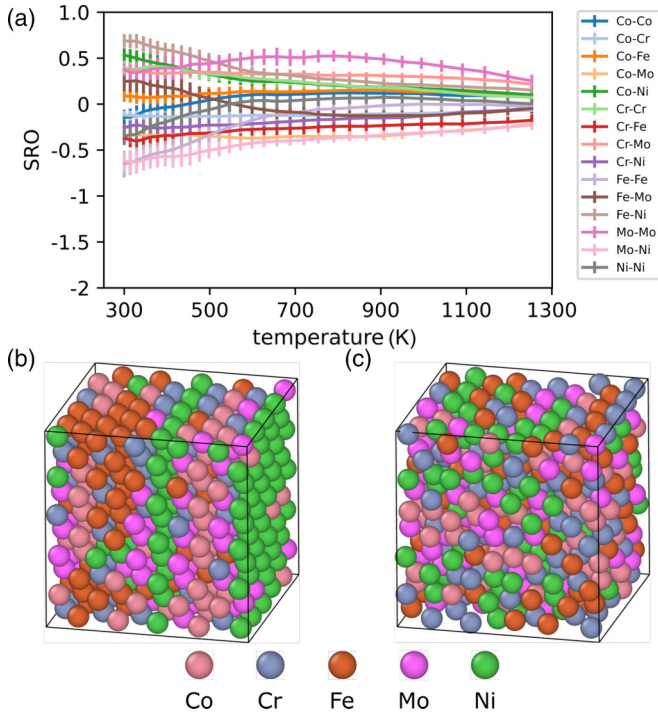


FIG. 13. (a) Cowley’s short-range (SRO) parameters for the first shell in CoCrFeMoNi HEA, shown for the ten replicas between 300 and 1253 K, averaged over the last 1000 steps and two independent runs. Good mixing of atomic species can be assumed due to the small values of SRO parameters. The y axis is adjusted to the example shown in Fig. 14 to facilitate comparison. (b), (c) Snapshot from MC/MD simulations at  $T = 300$  K and at  $T = 1253$  K, respectively. In the 300 K snapshot, two planes of Ni can be seen.

reduced further, a second transition occurs around 400 K. The most prominent structural transformation is the formation of (100) Ni planes, separated by (Co, Fe, Mn)-rich regions forming a layered superstructure. Figures 12(b) and 12(c) show snapshots of the simulations at 300 K and 720 K, that give an idea of the partially ordered structure of the two phases.

Substituting Mn with Mo changes the segregation behavior significantly [Fig. 13(a)]: the SRO parameters are generally smaller, with the largest segregation tendency found for the Mo-Ni atom pair. The tendency of Cr to form a cubic sublattice is less pronounced than CoCrFeMnNi, and one only sees the increase of SRO parameters at around 500 K. At low temperature, (100) planes of Ni form that are very similar to those observed in the Mn-based counterpart [Figs. 13(b) and 13(c)] that are separated by (Co, Fe, Mo)-rich regions. Given the sizable energy errors of the ML models, as well as those of the underlying DFT reference, one should not overinterpret the details of the structures we observe. Even if *fcc* CoCrFeMnNi is paramagnetic, neglect of magnetism in the presence of several elements which form ferromagnetic phases is worrisome (see, e.g., Ref. [97] for a thorough discussion of magnetism in CoCrFeMnNi and CoCrFeMoNi). That said, our observations provide strong indications of the tendency to form partly ordered phases with a complex structure, which, together with the low vacancy-mediated diffusivity [98] help explain the observed stability of HEAs that contain (Co, Fe, Cr, Ni). A

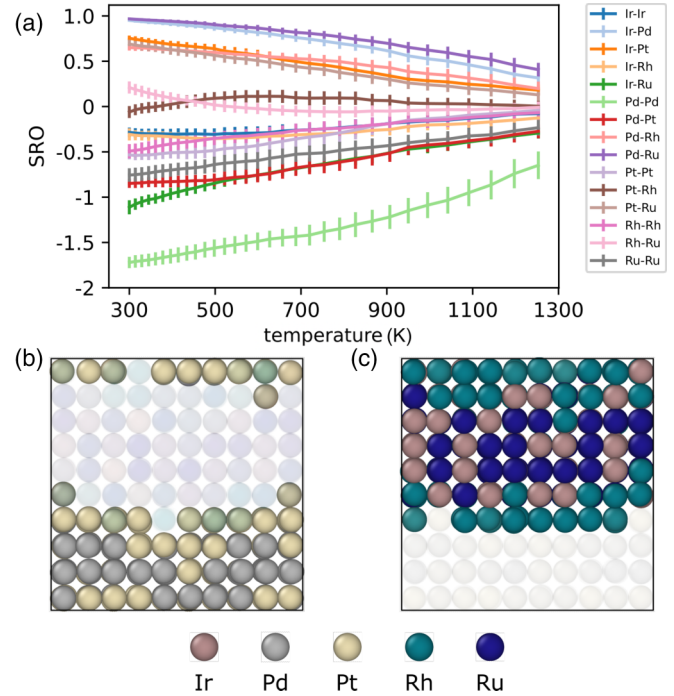


FIG. 14. (a) Cowley’s short-range parameters for the first shell in IrPdPtRhRu HEA, shown for the ten replicas between 500 and 933 K, averaged over the last 1000 frames and with an error estimation from independent repetition runs. The most pronounced local order can be seen for the Pd-Pd atom pair (light green line, mathematically smallest SRO). Demonstration of the phase segregation tendency by highlighting the (b) PdPt and (c) IrRhRu atoms in an MC/MD snapshot.

tendency to develop short-range ordering is consistent with previous simulations in other classes of HEAs [99] and with observation of phase separation in equimolar CoCrFeMnNi in high-mobility environments such as grain boundaries [100] or under deformation [101].

While the leading effect in CoCrFeMnNi and CoCrFeMoNi is the appearance of partial ordering at low temperatures, in the case of IrPdPtRhRu we observe clear-cut phase separation between a (Pd, Pt) and a (Ru, Ir, Rh) phase, with Rh accumulating preferentially at the interface between the two phases [see Figs. 14(b) and 14(c)]. The strong tendency to segregate is already evident in the high-temperature regime, where the system is visually well-mixed, but with large SRO parameters. This is in contrast to the experimental observation that this HEA forms a complex solid solution with random atom distribution [21,95]. As shown in the Supplemental Material [53], the large enthalpic gain arising from demixing is not an artifact of HEA25-4-NN, and the ML error on the free-energy change upon ordering is of the order of 3 meV/atom. These observations suggest that kinetic trapping, or finite-size effects associated with the synthesis in the form of nanoparticles, might be key to stabilize a homogeneous phase.

### VIII. CONCLUSIONS

The notion that different chemical elements may behave similarly when combined with others is one of the founding



principles of chemistry and is often used as guidance in the design of unique materials. We build a ML framework that incorporates this notion in the form of a linear compression of chemical space and succeed in training a potential that can describe with semiquantitative accuracy bulk phases of arbitrary combinations of 25 *d*-block elements. The physically motivated, intuitive functional form of the contraction allows us to analyze critically the model performance, allowing us to show that three to four dimensions suffice to capture the diversity of behavior of the transition metal block. The optimized values of the combination weights reveal relationships between the elements that match their arrangement in the periodic table to the point where we show it is possible to fill in the blanks for missing elements, with only a moderate loss of accuracy.

We use the potential to run an ambitious computational experiment, in which we attempt to equilibrate an equimolar mixture of all 25 elements, resulting in the formation of a disordered structure with strong element segregation. The affinity between elements is consistent with several known HEAs and allows us to define a data-driven version of the Hume-Rothery rules that could be further adapted to subsets of elements that are relevant for a given application. We also investigate in detail three specific compositions—the archetypal Cantor alloy CoCrFeMnNi, which we observe to be undergo a sequence of transitions towards complex ordered phases as the temperature is lowered; that arising from the Mn → Mo substitution, which also leads to similar, although less pronounced, ordering; and the noble metal alloy PdPtIrRuRh that shows a strong tendency to decompose into PdPt and IrRhRu phases.

We are only scratching the surface of what can be achieved within this framework. Extending the data set to an even

more diverse palette of compounds, and to structures that include molten and defective configurations, is an obvious direction for further improvements. A more systematic exploration of the design space of chemical compression is another promising research direction, even though doing so may sacrifice, at least in part, the interpretability of the linear contraction we use here. On a more application-focused front, a systematic study of the stability of four- and five-element HEAs along the same lines of the simulations of those which we present here, based on the current HEA25-4-NN model, will provide much-needed insights into the stability range of multi-principal-component alloys, guiding synthetic efforts towards compositions that are stable towards phase separation.

All data and code used to train the HEA25-4-NN model as well as the fitted parameters and code to run the simulations discussed in this paper are available in the Supplemental Material [53] or from publicly accessible repositories [102,103].

## ACKNOWLEDGMENTS

We thank William Curtin and Binglun Yin for sharing a data set we used in the early developments of this work, and Anirudh Natarajan for discussion on an early version of the paper. M.C. and N.L. acknowledge support from the NCCR MARVEL, funded by the Swiss National Science Foundation (Grant No. 182892), and from an industrial grant from BASF. G.F. acknowledges support by the Swiss Platform for Advanced Scientific Computing (PASC). Electronic-structure calculations were performed within the scope of a CSCS project (Project No. s1092).

- 
- [1] J.-W. Yeh, S.-K. Chen, S.-J. Lin, J.-Y. Gan, T.-S. Chin, T.-T. Shun, C.-H. Tsau, and S.-Y. Chang, *Adv. Eng. Mater.* **6**, 299 (2004).
- [2] B. Cantor, I. Chang, P. Knight, and A. Vincent, *Mater. Sci. Eng.: A* **375-377**, 213 (2004).
- [3] B. Cantor, *Prog. Mater. Sci.* **120**, 100754 (2021).
- [4] E. George, W. Curtin, and C. Tasan, *Acta Mater.* **188**, 435 (2020).
- [5] W. Li, D. Xie, D. Li, Y. Zhang, Y. Gao, and P. K. Liaw, *Prog. Mater. Sci.* **118**, 100777 (2021).
- [6] Y. Sun and S. Dai, *Sci. Adv.* **7**, eabg1600 (2021).
- [7] B. Wang, Y. Yao, X. Yu, C. Wang, C. Wu, and Z. Zou, *J. Mater. Chem. A* **9**, 19410 (2021).
- [8] W.-Y. Huo, S.-Q. Wang, W.-H. Zhu, Z.-L. Zhang, F. Fang, Z.-H. Xie, and J.-Q. Jiang, *Tungsten* **3**, 161 (2021).
- [9] Y. Zhang, D. Wang, and S. Wang, *Small* **18**, 2104339 (2022).
- [10] X. Huo, H. Yu, B. Xing, X. Zuo, and N. Zhang, *Chem. Rec.* **22**, e202200175 (2022).
- [11] G. Zhang, K. Ming, J. Kang, Q. Huang, Z. Zhang, X. Zheng, and X. Bi, *Electrochim. Acta* **279**, 19 (2018).
- [12] M. Bondesgaard, N. L. N. Broge, A. Mamakhel, M. Bremholm, and B. B. Iversen, *Adv. Funct. Mater.* **29**, 1905933 (2019).
- [13] M. W. Glasscott, A. D. Pendergast, S. Goines, A. R. Bishop, A. T. Hoang, C. Renault, and J. E. Dick, *Nat. Commun.* **10**, 2650 (2019).
- [14] Z. Jin, J. Lv, H. Jia, W. Liu, H. Li, Z. Chen, X. Lin, G. Xie, X. Liu, S. Sun, and H.-J. Qiu, *Small* **15**, 1904180 (2019).
- [15] S. D. Lacey, Q. Dong, Z. Huang, J. Luo, H. Xie, Z. Lin, D. J. Kirsch, V. Vattipalli, C. Pavinelli, W. Fan, R. Shahbazian-Yassar, D. Wang, and L. Hu, *Nano Lett.* **19**, 5149 (2019).
- [16] M. Liu, Z. Zhang, F. Okejiri, S. Yang, S. Zhou, and S. Dai, *Advanced Materials Interfaces* **6**, 1900015 (2019).
- [17] H.-J. Qiu, G. Fang, J. Gao, Y. Wen, J. Lv, H. Li, G. Xie, X. Liu, and S. Sun, *ACS Mater. Lett.* **1**, 526 (2019).
- [18] H.-J. Qiu, G. Fang, Y. Wen, P. Liu, G. Xie, X. Liu, and S. Sun, *J. Mater. Chem. A* **7**, 6499 (2019).
- [19] S. Gao, S. Hao, Z. Huang, Y. Yuan, S. Han, L. Lei, X. Zhang, R. Shahbazian-Yassar, and J. Lu, *Nat. Commun.* **11**, 2016 (2020).
- [20] K. Huang, B. Zhang, J. Wu, T. Zhang, D. Peng, X. Cao, Z. Zhang, Z. Li, and Y. Huang, *J. Mater. Chem. A* **8**, 11938 (2020).
- [21] D. Wu, K. Kusada, T. Yamamoto, T. Toriyama, S. Matsumura, I. Gueye, O. Seo, J. Kim, S. Hiroi, O. Sakata, S. Kawaguchi, Y. Kubota, and H. Kitagawa, *Chem. Sci.* **11**, 12731 (2020).

- [22] X. Chen, C. Si, Y. Gao, J. Frenzel, J. Sun, G. Eggeler, and Z. Zhang, *J. Power Sources* **273**, 324 (2015).
- [23] T. Löffler, H. Meyer, A. Savan, P. Wilde, A. G. Manjón, Y.-T. Chen, E. Ventosa, C. Scheu, A. Ludwig, and W. Schuhmann, *Adv. Energy Mater.* **8**, 1802269 (2018).
- [24] S. Li, X. Tang, H. Jia, H. Li, G. Xie, X. Liu, X. Lin, and H.-J. Qiu, *J. Catal.* **383**, 164 (2020).
- [25] J. Barranco and A. Pierna, *J. Non-Cryst. Solids* **354**, 5153 (2008).
- [26] C.-F. Tsai, K.-Y. Yeh, P.-W. Wu, Y.-F. Hsieh, and P. Lin, *J. Alloys Compd.* **478**, 868 (2009).
- [27] A.-L. Wang, H.-C. Wan, H. Xu, Y.-X. Tong, and G.-R. Li, *Electrochim. Acta* **127**, 448 (2014).
- [28] K. V. Yusenko, S. Riva, P. A. Carvalho, M. V. Yusenko, S. Arnaboldi, A. S. Sukhikh, M. Hanfland, and S. A. Gromilov, *Scr. Mater.* **138**, 22 (2017).
- [29] J.-W. Yeh, *JOM* **65**, 1759 (2013).
- [30] E. J. Pickering and N. G. Jones, *Int. Mater. Rev.* **61**, 183 (2016).
- [31] J. Sanchez, F. Ducastelle, and D. Gratias, *Physica A* **128**, 334 (1984).
- [32] C. M. Nataraj, A. van de Walle, and A. Samanta, *J. Phase Equilib. Diffus.* **42**, 571 (2021).
- [33] Y. Rao and W. Curtin, *Acta Mater.* **226**, 117621 (2022).
- [34] D. Ma, B. Grabowski, F. Körmann, J. Neugebauer, and D. Raabe, *Acta Mater.* **100**, 90 (2015).
- [35] D. Farkas and A. Caro, *J. Mater. Res.* **33**, 3218 (2018).
- [36] J. Byggmästar, K. Nordlund, and F. Djurabekova, *Phys. Rev. B* **104**, 104101 (2021).
- [37] C. W. Rosenbrock, K. Gubaev, A. V. Shapeev, L. B. Pártay, N. Bernstein, G. Csányi, and G. L. W. Hart, *npj Comput Mater* **7**, 24 (2021).
- [38] Y. Zhou, P. Srinivasan, F. Körmann, B. Grabowski, R. Smith, P. Goddard, and A. I. Duff, *Phys. Rev. B* **105**, 214302 (2022).
- [39] M. J. Willatt, F. Musil, and M. Ceriotti, *Phys. Chem. Chem. Phys.* **20**, 29661 (2018).
- [40] M. J. Willatt, F. Musil, and M. Ceriotti, *J. Chem. Phys.* **150**, 154110 (2019).
- [41] A. V. Shapeev, *Multiscale Model. Simul.* **14**, 1153 (2016).
- [42] R. Drautz, *Phys. Rev. B* **99**, 014104 (2019).
- [43] F. Musil, A. Grisafi, A. P. Bartók, C. Ortner, G. Csányi, and M. Ceriotti, *Chem. Rev.* **121**, 9759 (2021).
- [44] A. P. Bartók, R. Kondor, and G. Csányi, *Phys. Rev. B* **87**, 184115 (2013).
- [45] J. Behler, *J. Chem. Phys.* **134**, 074106 (2011).
- [46] G. C. Schatz, *J. Phys. Chem. Lett.* **10**, 5956 (2019).
- [47] N. Artrith, A. Urban, and G. Ceder, *Phys. Rev. B* **96**, 014112 (2017).
- [48] M. Gastegger, L. Schwiedrzik, M. Bittermann, F. Berzsenyi, and P. Marquetand, *J. Chem. Phys.* **148**, 241709 (2018).
- [49] J. P. Darby, J. R. Kermode, and G. Csányi, *npj Comput Mater* **8**, 166 (2022).
- [50] A. Goscinski, F. Musil, S. Pozdnyakov, J. Nigam, and M. Ceriotti, *J. Chem. Phys.* **155**, 104106 (2021).
- [51] J. Nigam, S. Pozdnyakov, and M. Ceriotti, *J. Chem. Phys.* **153**, 121101 (2020).
- [52] A. Paszke, S. Gross, F. Massa, A. Lerer, J. Bradbury, G. Chanan, T. Killeen, Z. Lin, N. Gimelshein, L. Antiga, A. Desmaison, A. Kopf, E. Yang, Z. DeVito, M. Raison, A. Tejani, S. Chilamkurthy, B. Steiner, L. Fang, J. Bai *et al.*, in *Advances in Neural Information Processing Systems 32*, edited by H. Wallach, H. Larochelle, A. Beygelzimer, F. dAlché-Buc, E. Fox, and R. Garnett (Curran Associates, Inc., Red Hook, New York, 2019), pp. 8024–8035.
- [53] See Supplemental Material at <http://link.aps.org/supplemental/10.1103/PhysRevMaterials.7.045802> for further details on the models, reference data, and example inputs of simulations.
- [54] G. Kresse and J. Furthmüller, *Phys. Rev. B* **54**, 11169 (1996).
- [55] G. I. Csonka, J. P. Perdew, A. Ruzsinszky, P. H. T. Philipsen, S. Lebègue, J. Paier, O. A. Vydrov, and J. G. Ángyán, *Phys. Rev. B* **79**, 155107 (2009).
- [56] G. Kresse and D. Joubert, *Phys. Rev. B* **59**, 1758 (1999).
- [57] H. J. Monkhorst and J. D. Pack, *Phys. Rev. B* **13**, 5188 (1976).
- [58] I. Novikov, B. Grabowski, F. Körmann, and A. Shapeev, *npj Comput Mater* **8**, 13 (2022).
- [59] C. S. Wang, R. E. Prange, and V. Korenman, *Phys. Rev. B* **25**, 5766 (1982).
- [60] J. Kubler, K. H. Hock, J. Sticht, and A. R. Williams, *J. Phys. F: Met. Phys.* **18**, 469 (1988).
- [61] H. J. Kulik, M. Cococcioni, D. A. Scherlis, and N. Marzari, *Phys. Rev. Lett.* **97**, 103001 (2006).
- [62] Z. Li, A. Ludwig, A. Savan, H. Springer, and D. Raabe, *J. Mater. Res.* **33**, 3156 (2018).
- [63] G. Imbalzano, A. Anelli, D. Giofré, S. Klees, J. Behler, and M. Ceriotti, *J. Chem. Phys.* **148**, 241730 (2018).
- [64] S. Haykin, *Neural Networks: A Comprehensive Foundation* (Prentice Hall PTR, 1994).
- [65] F. Musil, M. Veit, A. Goscinski, G. Fraux, M. J. Willatt, M. Stricker, and M. Ceriotti, *J. Chem. Phys.* **154**, 114109 (2021).
- [66] J. Behler and M. Parrinello, *Phys. Rev. Lett.* **98**, 146401 (2007).
- [67] D. C. Liu and J. Nocedal, *Math. Program.* **45**, 503 (1989).
- [68] H. C. Andersen, *J. Chem. Phys.* **72**, 2384 (1980).
- [69] G. Bussi, T. Zykova-Timan, and M. Parrinello, *J. Chem. Phys.* **130**, 074101 (2009).
- [70] M. Ceriotti, G. Bussi, and M. Parrinello, *J. Chem. Theory Comput.* **6**, 1170 (2010).
- [71] G. Bussi, D. Donadio, and M. Parrinello, *J. Chem. Phys.* **126**, 014101 (2007).
- [72] N. Metropolis, A. W. Rosenbluth, M. N. Rosenbluth, A. H. Teller, and E. Teller, *J. Chem. Phys.* **21**, 1087 (1953).
- [73] D. J. Earl and M. W. Deem, *Phys. Chem. Chem. Phys.* **7**, 3910 (2005).
- [74] V. Kapil, M. Rossi, O. Marsalek, R. Petraglia, Y. Litman, T. Spura, B. Cheng, A. Cuzzocrea, R. H. Meißner, D. M. Wilkins, B. A. Helfrecht, P. Juda, S. P. Bienvenue, W. Fang, J. Kessler, I. Poltavsky, S. Vandenbrande, J. Wieme, C. Corminboeuf, T. D. Kühne *et al.*, *Comput. Phys. Commun.* **236**, 214 (2019).
- [75] G. Imbalzano and M. Ceriotti, *Phys. Rev. Mater.* **5**, 063804 (2021).
- [76] R. Petraglia, A. Nicolai, M. D. Wodrich, M. Ceriotti, and C. Corminboeuf, *J. Comput. Chem.* **37**, 83 (2016).
- [77] F. A. Faber, A. Lindmaa, O. A. Von Lilienfeld, and R. Armiento, *Phys. Rev. Lett.* **117**, 135502 (2016).
- [78] V. L. Deringer and G. Csányi, *Phys. Rev. B* **95**, 094203 (2017).
- [79] V. L. Deringer, A. P. Bartók, N. Bernstein, D. M. Wilkins, M. Ceriotti, and G. Csányi, *Chem. Rev.* **121**, 10073 (2021).
- [80] G. T. Seaborg, Plutonium: The Ormery Element, *Chemistry*, **37**, 12 (1964).

- [81] D. Sheppard, G. Henkelman, and O. A. von Lilienfeld, *J. Chem. Phys.* **133**, 084104 (2010).
- [82] L. Bellaiche and D. Vanderbilt, *Phys. Rev. B* **61**, 7877 (2000).
- [83] C. Chen and S. P. Ong, *Nat. Comput. Sci.* **2**, 718 (2022).
- [84] A. Jain, S. P. Ong, G. Hautier, W. Chen, W. D. Richards, S. Dacek, S. Cholia, D. Gunter, D. Skinner, G. Ceder, and K. A. Persson, *APL Mater.* **1**, 011002 (2013).
- [85] J. M. Cowley, *Phys. Rev.* **77**, 669 (1950).
- [86] U. Mizutani, *MRS Bull.* **37**, 169 (2012).
- [87] F. Waag, Y. Li, A. R. Ziefuß, E. Bertin, M. Kamp, V. Duppel, G. Marzun, L. Kienle, S. Barcikowski, and B. Gökce, *RSC Adv.* **9**, 18547 (2019).
- [88] H. Peng, Y. Xie, Z. Xie, Y. Wu, W. Zhu, S. Liang, and L. Wang, *J. Mater. Chem. A* **8**, 18318 (2020).
- [89] J. Tang, J. Xu, Z. Ye, X. Li, and J. Luo, *J. Mater. Sci. Technol.* **79**, 171 (2021).
- [90] S. Schumacher, S. Baha, A. Savan, C. Andronescu, and A. Ludwig, *J. Mater. Chem. A* **10**, 9981 (2022).
- [91] G. Cui, B. Han, Y. Yang, Y. Wang, and H. Chunyang, *J. Mater. Res. Technol.* **9**, 2598 (2020).
- [92] Z. Wang, J. Jin, G.-H. Zhang, X.-H. Fan, and L. Zhang, *Corros. Sci.* **208**, 110661 (2022).
- [93] T. A. Batchelor, J. K. Pedersen, S. H. Winther, I. E. Castelli, K. W. Jacobsen, and J. Rossmeisl, *Joule* **3**, 834 (2019).
- [94] N. L. N. Broge, M. Bondesgaard, F. Søndergaard-Pedersen, M. Roelsgaard, and B. B. Iversen, *Angew. Chem.* **132**, 22104 (2020).
- [95] J. K. Pedersen, C. M. Clausen, O. A. Krysiak, B. Xiao, T. A. A. Batchelor, T. Löffler, V. A. Mints, L. Banko, M. Arenz, A. Savan, W. Schuhmann, A. Ludwig, and J. Rossmeisl, *Angew. Chem.* **133**, 24346 (2021).
- [96] G. Lee, N.-A. Nguyen, V.-T. Nguyen, L. L. Larina, E. Chuluunbat, E. Park, J. Kim, H.-S. Choi, and M. Keidar, *J. Solid State Chem.* **314**, 123388 (2022).
- [97] J. Šebesta, K. Carva, and D. Legut, *Phys. Rev. Mater.* **3**, 124410 (2019).
- [98] K.-Y. Tsai, M.-H. Tsai, and J.-W. Yeh, *Acta Mater.* **61**, 4887 (2013).
- [99] S. Chen, Z. H. Aitken, S. Pattamatta, Z. Wu, Z. G. Yu, D. J. Srolovitz, P. K. Liaw, and Y.-W. Zhang, *Nat. Commun.* **12**, 4953 (2021).
- [100] M. Glienke, M. Vaidya, K. Gururaj, L. Daum, B. Tas, L. Rogal, K. Pradeep, S. V. Divinski, and G. Wilde, *Acta Mater.* **195**, 304 (2020).
- [101] K. H. Lee, S.-K. Hong, and S. I. Hong, *Materialia* **8**, 100445 (2019).
- [102] <https://zenodo.org/record/7836106>.
- [103] <https://archive.materialscloud.org/record/2023.57>.