



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

Predicting activation energies for vacancy-mediated diffusion in alloys using a transition-state cluster expansion

Chenyang Li, Thomas Nilson, Liang Cao, and Tim Mueller

Phys. Rev. Materials **5**, 013803 — Published 25 January 2021

DOI: [10.1103/PhysRevMaterials.5.013803](https://doi.org/10.1103/PhysRevMaterials.5.013803)

Predicting activation energies for vacancy-mediated diffusion in alloys using a transition-state cluster expansion

Chenyang Li,¹ Thomas Nilson,¹ Liang Cao,¹ and Tim Mueller^{1,*}

¹Department of Materials Science and Engineering, Johns Hopkins University, Baltimore, MD 21218, USA

* Email: tmueller@jhu.edu

ABSTRACT

Kinetic Monte Carlo models parameterized by first principles calculations are widely used to simulate atomic diffusion. However, accurately predicting the activation energies for diffusion in substitutional alloys remains challenging due to the wide variety of local environments that may exist around the diffusing atom. We address this challenge using a cluster expansion model that explicitly includes a sublattice of sites representing transition states and assess its accuracy in comparison with other models, such as the broken bond model and a model related to Marcus theory, by modeling vacancy-mediated diffusion in Pt-Ni nanoparticles. We find that the prediction error of the cluster expansion is similar to that of other models for small training sets, but with larger training sets the cluster expansion has a significantly lower prediction error than the other models with comparable execution speed. Of the simpler models, the model related to Marcus theory yields predictions of nanoparticle evolution that are most similar to those of the cluster expansion, and a weighted average of the two approaches has the lowest prediction error for activation energies across all training set sizes.

I. INTRODUCTION

Atomic diffusion plays an important role in determining various material properties such as the structure and activity of nanocatalysts [1-5], the conductivity and stability of batteries [6-9] and fuel cells [10-13], and the electrical properties of semiconductor devices [14-17]. Atomic diffusion in materials often consists of hops between local minima on the potential energy surface, with atoms vibrating about these local minima in the time between hops. If the rates of these transitions are known, the simulation of diffusion can be significantly accelerated through the use of kinetic Monte Carlo (KMC) [18,19], which simulates only the discrete jumps between local minima rather than continuous time evolution of the system. The key to a KMC simulation is the rapid calculation of transition rates between local minima. By transition state theory [20], these rates depend exponentially on the activation energies of the hops, making the rapid and accurate calculation of activation energies critically important to the success of a KMC simulation.

Several methods for predicting activation energies in KMC simulations have been proposed and explored in the literature. These include the broken bond model [21-28] and a model that is similar to Marcus theory [29-31], which have been used in KMC simulations to study diffusion-related properties of materials. These methods are relatively simple and can be applied to a variety of material systems. However, the simplicity of these models limits their accuracy, and there is no way to systematically improve them if the predicted activation energies are not sufficiently accurate. An appealing alternative is to use systematically improvable machine-learned energy models such as cluster expansions [32-34] and interatomic potentials [35-42] for the rapid and accurate prediction of activation energies. As the cluster expansion is a discrete model explicitly designed to calculate the energies of local minima on the potential energy surface, it is particularly well suited for KMC.

The cluster expansion has become a valuable tool for studying atomic order and structure-property relationships in alloys [43-58], but it is rarely used to model diffusion [33,59,60]. Soisson and Fu [61,62], building off of work by Bouar and Soisson [63], developed a cluster-expansion-like lattice model to predict activation energies in which the energy of the transition state is a sum of nearest-neighbor pair interactions. In 2001, Van der Ven *et al.* demonstrated how cluster expansions could be used to predict activation energies in bulk Li_xCoO_2 by constructing two cluster expansions: a global cluster expansion of the local minima on the potential energy surface, and a local cluster expansion around the hopping atom that calculated a “kinetically resolved activation barrier” [33]. The combination of the two cluster expansions enabled the rapid calculation of activation energies. In a separate paper, Van der Ven *et al.* used a cluster expansion to predict the energies of both end states and local minima along the diffusion path [64]. To calculate the energies of transition states between local minima, they added a constant value to the higher-energy local minimum, where the constant value depended on whether or not the higher-energy site shared a face with a nearby occupied site. In 2010, Bhattacharya and Van der Ven adapted this approach for situations in which an atom at the intermediate site along the diffusion path is at a saddle point rather than a local minimum [65]. They accomplished this by using nudged elastic band method [66,67] to calculate the energies of structures in the training set in which the atom was at the intermediate site. This approach enabled them to calculate the energies of both the initial and transition states using a single global many-body cluster expansion.

Here, we present a method equivalent to this latter approach, in which transition states are explicitly included in a global cluster expansion as a set of sublattice sites, for modeling vacancy-mediated diffusion in Pt-Ni nanoparticles. The advantage to using a single global cluster expansion, rather than separate global and local cluster expansions, is that a single cluster expansion may be

more compatible with general-purpose cluster expansion software packages and easier to extend to include correlated hops in which multiple atoms change sites at the same time. Pt-Ni nanoparticles have been widely studied as potential catalysts for the oxygen reduction reaction [1,68-71] and we present a challenging system for predicting activation energies due to the variety of local environments that exist both in the bulk and on the surface. On this system we find the cluster expansion approach yields significantly more accurate activation energies than alternative leading approaches with little to no increase in computational cost for the kinetic Monte Carlo simulation.

II. METHODS

A. DFT

Density functional theory (DFT) [72] calculations were performed using the Vienna Ab initio Simulation Package (VASP) [73] with the revised Perdew-Burke-Ernzerhof (RPBE) [74,75] exchange-correlation functional. The Pt_pv_GW and Ni PBE projector-augmented wave (PAW) [76] potentials were used. A single k -point at the center of the Brillouin zone was used for each nanoparticle. For bulk materials, the Brillouin zone was sampled using generalized Monkhorst-Pack grids generated by the k -point grid server [77] with a minimum distance of 46.5 Å between real space lattice points. Second-order Methfessel-Paxon smearing [78] with a width of 0.2 eV was used to set partial occupancies. The convergence criteria for the electronic self-consistency iteration and the ionic relaxation were set to be 10^{-4} eV and 0.03 eV/Å, respectively. The climbing image nudged elastic band (CI-NEB) [66,67] method was used to calculate the activation energies for atomic diffusion, and the calculations were considered to be sufficiently converged when the maximum force component perpendicular to the diffusion path was below 0.05 eV/Å.

B. Cluster expansion

Cluster expansions are generalized Ising models that account for many-body interactions [32,79] and are commonly used to study atomic order in substitutional materials [47,54-56,80-82]. In this application, the arrangement of atoms in the material is expressed as a set of discrete sites that are occupied by different elements. In the cluster expansion constructed here, the sites are arranged on an fcc lattice and each site is occupied by a Pt atom, Ni atom, or vacancy. We fit the parameters of the cluster expansion, known as effective cluster interactions (ECI), to DFT-calculated energies using the Bayesian approach [83], which has been shown to effectively improve the accuracy of the cluster expansion for a given training set size [54,81,83]. Specifically, we assumed exponential decay in the width of the prior probability distribution with respect to the number of sites in each cluster and the maximum distance between sites.

To calculate activation energies for diffusion using the cluster expansion we added a sublattice of “transition state” sites halfway between nearest-neighbor sites (Figure 1a). Each of these transition state sites corresponds to a saddle point on the potential energy surface for a hop between neighboring sites. During a single diffusive hop, an atom starts at one of the fcc sites, moves to an adjacent transition state site, and then ends up in a nearest-neighbor fcc site (Figure 1b). By comparing the calculated energy of the system when the atom is at the transition state site with the energy of the system when the atom is in the initial fcc site, we directly predict the activation energy from the cluster expansion. To train the cluster expansion, the energies of the configurations where an atom is at a transition state site were obtained from DFT nudged elastic band calculations.




Figure 1. (a) A schematic of transition state cluster expansion. Large gray and brown spheres represent Pt and Ni, respectively, on lattice sites (local minima on the potential energy surface). Small green dots represent saddle points for hops between two local minima. For illustration purposes a two-dimensional lattice is shown, but the model in this paper is built on a three-dimensional fcc lattice. (b) A set of three training structures representing a diffusive hop, with the diffusing species at the transition state site circled in green. From left to right: initial, transition, and final state.

We constrained the set of allowed configurations to prevent simultaneous occupation of a transition state site and an adjacent fcc site, consistent with the mechanism of a vacancy-mediated hop. We also included a constraint that prevented two transition state sites from being occupied simultaneously, effectively allowing only a single atom to move during any hop. We note that this framework could be extended in future work by removing this second constraint. This would allow for collective diffusion, in which multiple atoms hop simultaneously, which is particularly likely to occur on surfaces [84-86].

The constraints on the set of allowed configurations make the cluster expansion overdetermined, i.e. there are more possible cluster functions than possible configurations. The cluster expansion can therefore be simplified by removing unnecessary functions. Here we accomplish this by

removing all functions associated with clusters that include two sites that cannot be simultaneously occupied. As the number of remaining cluster functions is then equal to the number of allowed configurations, and these functions are linearly independent, they form a complete basis for the constrained configuration space. The proofs of these statements are provided in the Appendix. We note that this simplification procedure is not limited to studies of diffusion and can be used in any cluster expansion in which some sites cannot be simultaneously occupied.

A training set of 299 relaxed structures was generated using DFT for fitting the cluster expansion. The training set contained five bulk structures and the initial, transition, and final states of diffusive hops in different nanoparticles (Figure 1b). Each of the five bulk structures was included twice to ensure accuracy in the bulk limit. The cluster expansion was truncated to include the empty cluster, point clusters, all two-body clusters up to a cutoff distance of 8 Å, all three-body clusters up to a cutoff distance of 6 Å, all four-body clusters up to a cutoff distance of 4 Å, all five-body clusters up to a cutoff distance of 4 Å, and all six-body clusters up to a cutoff distance of 4 Å. The resulting cluster expansion contained 1097 symmetrically distinct cluster functions. In addition to the 299 energies in the training set, we also included 196 activation energies (energy differences between the transition state and the two corresponding end states) explicitly in the fitting to improve accuracy, as the expressions for these energies only include the terms that change between the initial and transition states. An additional regularization parameter was introduced to distinguish the expected magnitude of ECI for clusters that include and do not include transition states. Additional information about how the cluster expansion was fit can be found in section 1 of the Supplementary Material [87].

III. RESULTS AND DISCUSSION

We begin by assessing the predictive accuracy of the transition-state cluster expansion on both formation energies and activation energies. The root-mean-square leave-one-out cross validation (LOOCV) errors are 1.203 meV/atom for the formation energies and 0.127 eV per particle for the activation energies (Figure 2), relative to DFT calculations. The LOOCV error of the activation energies is equivalent to 0.944 meV/atom, which is of the same order of magnitude as the error for the formation energies. This LOOCV error compares favorably to validation errors for other machine learning methods for predicting activation energies (Table S1 [35-42]), especially considering that the DFT data set contains hops in a wide variety of coordination environments, including both in the bulk and on the surface.

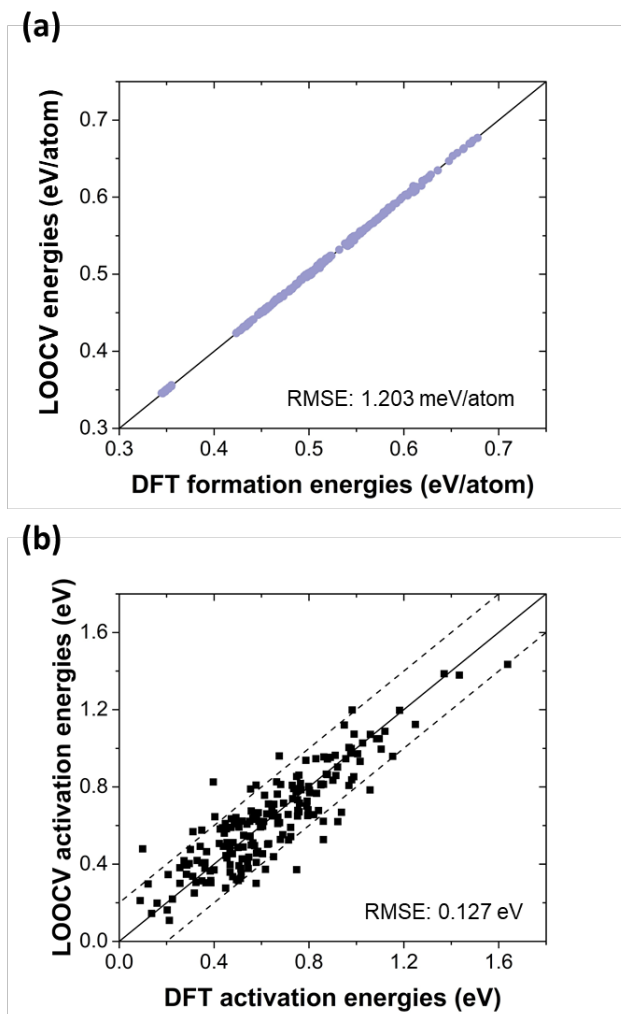


Figure 2. Leave-one-out cross validation (LOOCV) of a) the formation energies and b) activation energies from the transition-state cluster expansion. The dashed lines in (b) indicate ± 0.2 eV deviation from perfect agreement between DFT-calculated and cluster-expansion-predicted activation energies.

To further validate the accuracy of the transition-state cluster expansion for predicting the energies of local minima on the potential energy surface, we have used it to predict the equilibrium surface composition profile of a 4.5 nm cuboctahedral Pt_3Ni nanoparticle ($\text{Pt}_{2547}\text{Ni}_{849}$) at 333 K (Figure S2, Supplementary Material section 2.1 [87]), using Metropolis Monte Carlo [88] simulations in a canonical ensemble. The near-surface composition profile (Figure S2) is similar to those which have been experimentally [68] and computationally [69,89] determined for an extended $\text{Pt}_3\text{Ni}(111)$

surface, consistent with the fact that the surface of the cuboctahedral Pt_3Ni nanoparticle consists of mainly (111) facets.

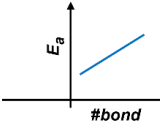
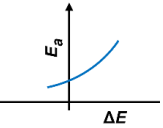
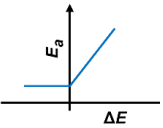
We have compared the accuracy of transition-state cluster expansion for predicting activation energies to three other simple models that have been developed for KMC simulations: the broken bond model [21], a model related to Marcus theory [29], and a model we have previously used in which the activation energy for a hop from a higher-energy state to a lower-energy state (i.e. a “downhill” hop) is a constant value [1,90]. Each of these models was trained using the same data used to train the cluster expansion. In each of these models other than the broken bond model, the cluster expansion is used to calculate the energy of each end state and the transition state energy is expressed as a function of the end state energies. The functional forms of the simple models we have evaluated are shown in Table 1.

In the broken bond model [21], the activation energy for diffusion is calculated as a linear function of the number of nearest-neighbor bonds of each type (e.g. Pt-Pt, Pt-Ni, or Ni-Ni) that are broken when the hopping atom leaves its initial state. One of the drawbacks of this model is that it does not satisfy detailed balance, as the difference in activation energies between the forward and reverse hops will generally not be the same as the difference in energies between the end states. We have also evaluated a broken-bond approach [24] in which detailed balance is restored by calculating a kinetically resolved activation (KRA) energy [33] as a function of the average number of bonds of each type at the initial and final states.

In the model related to Marcus theory [29], the potential energy surface of the initial and final state of the reaction as approximated using a simple parabolic form. In this model the activation energy can be analytically expressed as a function of the difference between the initial and final states,

ΔE (Table 1, Supplementary Material section 3 [87]). We will refer to this model as the “parabolic potential” model.

Table 1. Models used as comparisons to the cluster expansion in this work. The fitted parameters are determined by minimizing the RMS LOOCV error. n^i and n^f are the number of bonds at the initial and final states, respectively.

	Functional forms of activation energy	Fitted parameters	Schematics
Broken bond	$E_a = n_{Pt-Pt}^i E_{Pt-Pt} + n_{Pt-Ni}^i E_{Pt-Ni} + n_{Ni-Ni}^i E_{Ni-Ni}$	$E_{Pt-Pt} = 0.097$ eV $E_{Pt-Ni} = 0.066$ eV $E_{Ni-Ni} = 0.067$ eV	
Broken bond (KRA)	$E_a = \frac{1}{2} \Delta E + \frac{1}{2} (n_{Pt-Pt}^i + n_{Pt-Pt}^f) E_{Pt-Pt} + \frac{1}{2} (n_{Pt-Ni}^i + n_{Pt-Ni}^f) E_{Pt-Ni} + \frac{1}{2} (n_{Ni-Ni}^i + n_{Ni-Ni}^f) E_{Ni-Ni}$	$E_{Pt-Pt} = 0.097$ eV $E_{Pt-Ni} = 0.067$ eV $E_{Ni-Ni} = 0.066$ eV	
Parabolic potential	$E_a = \begin{cases} 0, & \Delta E < -4b \\ \frac{(\Delta E + 4b)^2}{16b}, & -4b \leq \Delta E \leq 4b \\ \Delta E, & \Delta E > 4b \end{cases}$	$b = 0.628$ eV	
Constant activation energy	$E_a = \begin{cases} c, & \Delta E < 0 \\ c + \Delta E, & \Delta E \geq 0 \end{cases}$	$c = 0.510$ eV	
Weighted average	$E_a = f E_a^{\text{CE}} + (1 - f) E_a^{\text{parabolic}}$	$f = 0.75$	

The final model, which we have used in previous work [1,90], uses a constant activation energy for hops to a lower-energy state and the same constant energy plus ΔE otherwise. The advantage

of this model is that the value of the constant activation energy only affects the prefactor of the KMC simulation (Supplementary Material section 2.3 [87]), so it is possible to simulate the dynamics of diffusion without including any NEB-calculated activation energies in the training data. We will refer to this model as the “constant activation energy” model.

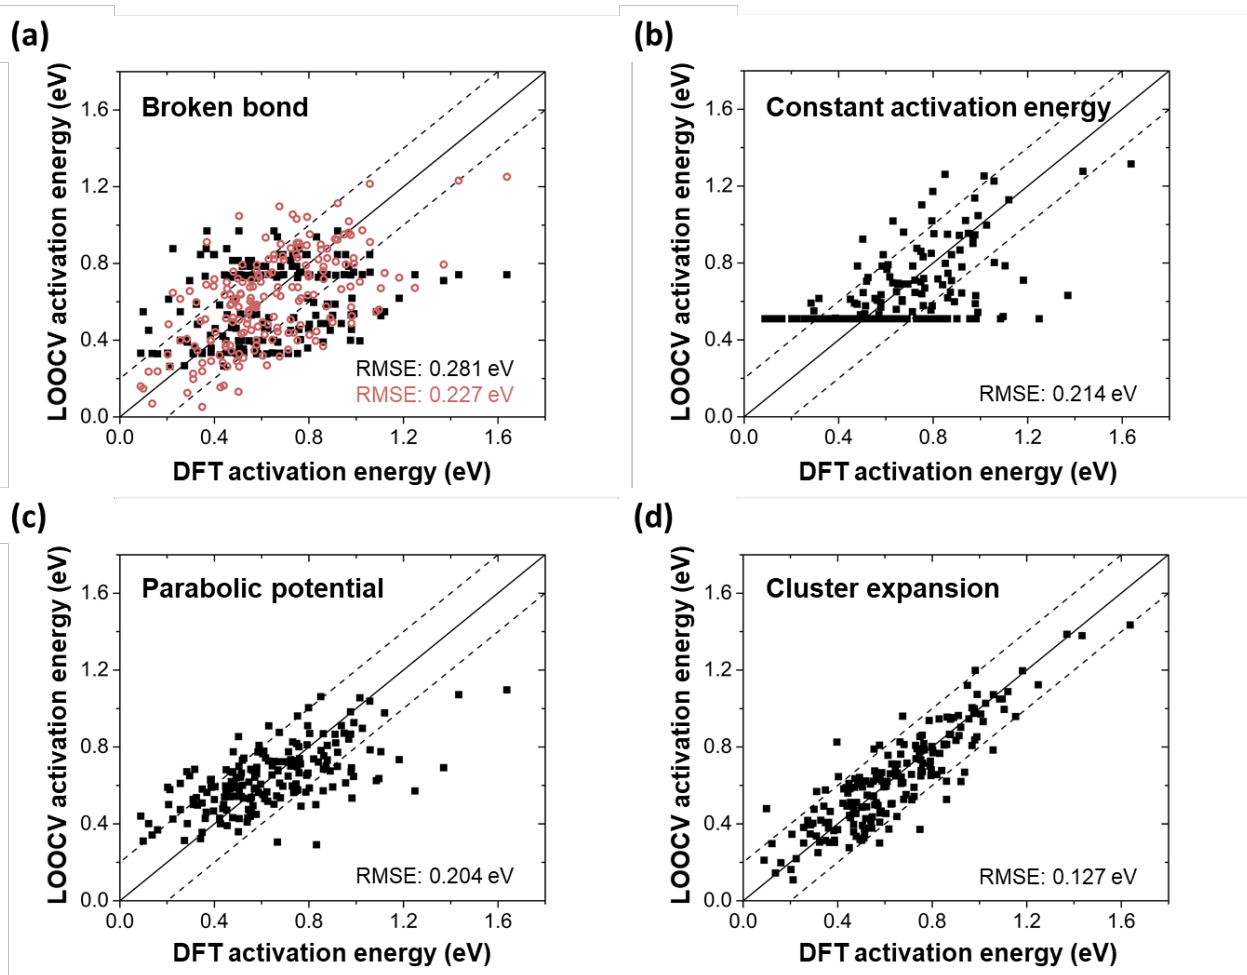


Figure 3. Leave-one-out cross validation (LOOCV) of the activation energy from the four methods against the known DFT activation energy in the training set. In a) the black and red data points are the simple broken bond and the KRA broken bond model, respectively. The dashed lines indicate ± 0.2 eV deviation from perfect agreement.

The simple broken bond model has the largest root-mean-square LOOCV error (0.281 eV), followed by the KRA broken bond model (0.227 eV), the constant activation energy model (0.214

eV), the parabolic potential model (0.204 eV), and the cluster expansion (0.127 eV) (Figure 3)**Error! Reference source not found..** The relatively low error for the cluster expansion is due largely to the fact that it is systematically improvable. As a result, it includes far more parameters (1097) than the other models. Because the simple broken bond model only captures the nearest-neighbor environment around the initial state of the diffusing atom and its functional form establishes an upper bound for the activation energy (~ 1 eV), it predicts many hops with identical activation energies regardless of the end states and it underestimates the activation energies for many hops with large activation energies. The KRA broken bond model is significantly more accurate, likely because it does take into account the end states and makes the forward and reverse hops consistent. In particular, because this model depends explicitly on the energy difference between the initial and final states, it does not impose an upper bound on the activation energy as the broken bond model does. The constant activation energy model establishes a lower bound of the activation energy that is assigned to all downhill hops (Figure 3b). This causes overestimation and underestimation of small and large activation energies, respectively. The trend for the parabolic potential model (Figure 3c) is better, although it still significantly overestimates and underestimates some small and large activation energies respectively. A comparison of how the constant activation energy model and the parabolic potential models predict the activation energy as a function of reaction energy is provided in Figure S3.

It is often possible to create a linear combination of multiple models to create a new model with improved predictive accuracy, an approach generally known as ensemble learning [91]. To evaluate this approach, we created a model consisting of a weighted average of the cluster expansion and the parabolic potential model (Table 1). The relative weights of the two models were determined in a way to minimize the leave-one-out cross validation error (Table 1 and

Supplementary Material section 4 [87]). We note that the weights determined this way, 0.75 for the cluster expansion and 0.25 for the parabolic potential model, are similar to what we would have obtained using inverse variance weighting [92] (0.72 for the cluster expansion and 0.28 for the parabolic potential model), using the leave-one-out cross validation errors to estimate the variance. The newly constructed weighted average of the cluster expansion and parabolic potential models further lowers the LOOCV error by 10% (13 meV) relative to cluster expansion itself (Figure S4, Supplementary Material [87]). This can be rationalized by the fact that these two models have error distributions that are largely uncorrelated due to their very different formalisms (Figure S5).

The improved performance of the cluster expansion comes with additional computational cost required to generate the training data. To assess the trade-off between cost and performance, we have evaluated the accuracy of the different methods as a function of training set size by randomly partitioning our data set into training sets and test sets, with models trained on the training sets and evaluated using the corresponding test sets. We started by randomly selecting 20% of the total data set as test set, training each model on all the remaining data, and then testing each model on the test set. This procedure was repeated 10 times to calculate the average root-mean-square test error. We then repeated the procedure using randomly-selected training sets that were 80%, 60%, 40%, and 20% as large as the first training set. The averaged test set root-mean-square errors for the 10 different runs at each size are shown in Figure 4, with the standard deviations colored as shaded regions.




Figure 4. Test set root mean square (RMS) error as a function of training set size, expressed as a percentage of the remaining data set (excluding the test set), which are 234 structures. The test set contains 60 structures. The standard deviations of the test set errors are colored as shaded regions.

For the cluster expansion, the test set errors decrease significantly with increasing training set size, whereas for the broken bond models, constant activation energy model, and parabolic potential model the test set errors remain relatively constant (Figure 4). This can be understood by the more flexible form of the cluster expansion compared to the simpler models (Table 1). These results indicate that when little training data is available, the parabolic potential, constant activation energy, and KRA broken bond models can predict activation energies with accuracy comparable to the cluster expansion, but if higher accuracy is desired the cluster expansion can be improved through the generation of additional training data. The weighted average between the cluster expansion and parabolic potential model has the lowest error at every training set size and is an appealing option especially for relatively small training sets. For the system considered here, the weighted average model achieves the same accuracy of the cluster expansion with only 60% of the training data (Figure 4).

To investigate how the different models may affect the structural evolution of a Pt₃Ni nanoparticle and compare model execution speeds, we have performed KMC simulations on a truncated octahedral Pt₃Ni nanoparticle with randomly initialized atomic order (Pt₃₄₁₁Ni₁₁₃₇, with a diameter of approximately 6.2 nm) at 1000 K. The elevated temperature was chosen to accelerate dynamics and reduce the computational cost of the KMC simulations. In each KMC simulation, atoms were only allowed to hop into neighboring vacant sites and the activation energy was computed using different models. We also considered Ni dissolution, as significant Ni loss is observed experimentally [1,90,93]. We only allowed Ni dissolution from surface sites with a coordination number less than 10, and as long as all of the neighboring atoms were left with at least 3 nearest-neighbor atoms after Ni dissolution. For simplicity, for all models we assigned a zero activation energy for Ni dissolution, which generally occurs readily from the surface [1]. More details are provided in Supplementary Material section 2.2 [87]. The snapshots of the Pt-Ni nanoparticles after the KMC simulations are shown in Figure 5a-e.

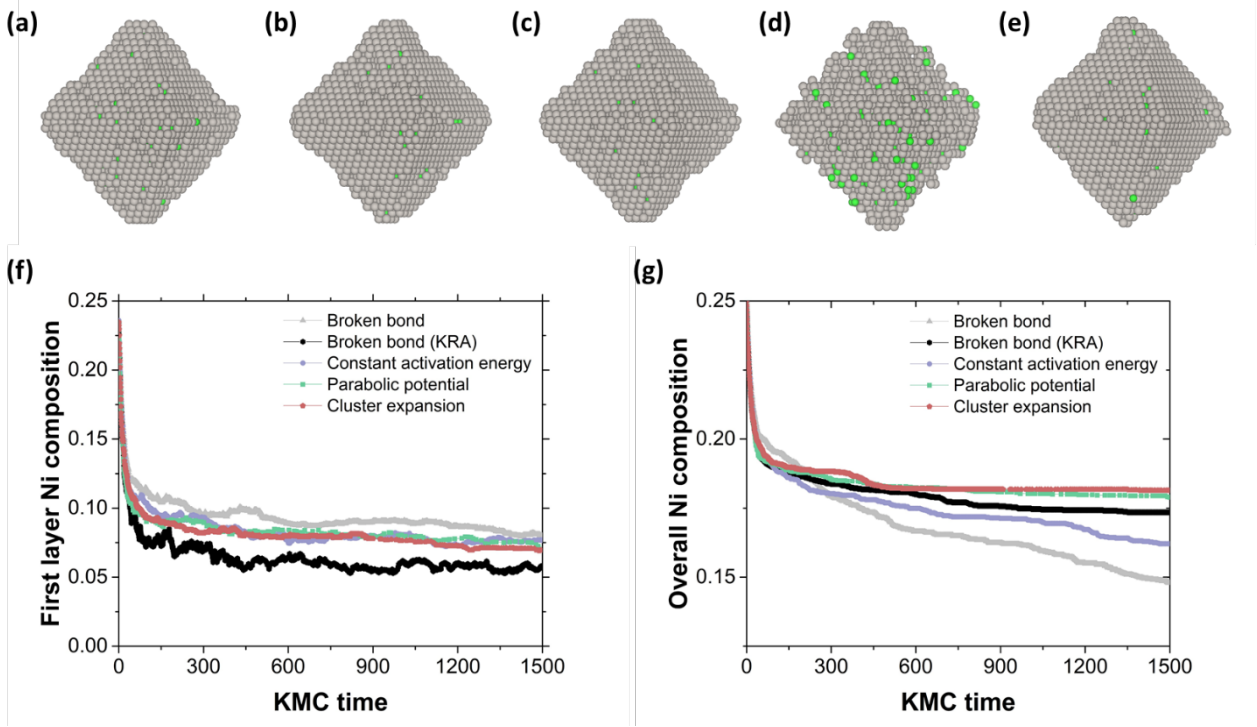


Figure 5. (a-e) Snapshots of Pt-Ni nanoparticles after the KMC simulations at 1000 K using the a) cluster expansion, b) parabolic potential model, c) constant activation energy model, d) simple broken bond model, and e) KRA broken bond model. Grey spheres represent Pt and green spheres represent Ni. (f, g) The first layer Ni composition and overall Ni composition from the KMC simulations as a function of KMC time.

The evolution of Ni composition in the first (outermost) layer is largely similar for the cluster expansion, parabolic potential model, and constant activation energy model (Figure 5f), where a Pt-rich shell is formed, consistent with our previous KMC work and experimental work [1,90]. The simple broken bond model produced a higher Ni composition in the first layer than that of the other models, but with the lowest overall Ni composition (Figure 5g), indicating that the Ni atoms are exchanged more rapidly from inner layers to the surface of the nanoparticle. The KRA broken bond model, however, shows completely different kinetics from the simple broken bond model, with a faster and more complete formation of a Pt shell and Pt(111) surface. This is likely due to

the fact that the KRA broken bond model satisfies detailed balance, resulting in evolution towards a lower energy state.

The overall Ni dissolution rate (Figure 5g) largely follows the trend of the prediction errors of the five models. The simple broken bond model has the largest error and the cluster expansion has the smallest, with the parabolic potential and the constant activation energy model in between. The large prediction error of the simple broken bond model is responsible for promoting more hops with unfavorable change in energy, thus creating more defects and a more porous structure (Figure 5d). The less accurate models, such as the broken bond models and constant activation energy model, have less spread in the predicted activation energies and thus might artificially accelerate rare events that could lead Ni to segregate to the surface and dissolve, resulting in lower overall Ni compositions (Figure 5g).

Although the cluster expansion is much more accurate for predicting activation energies, it does not significantly affect the shape evolution of the Pt₃Ni nanoparticle compared to the parabolic potential model or the constant activation energy model (Figure 5a-c, Figure 5f), except for the amount of predicted Ni loss. We analyze the nearest-neighbor environments of the Pt-Ni nanoparticles from the five models at the same overall Ni composition (18% Ni, Supplementary Material section 2.2 [87]). The values of N_{Pt-Pt} , N_{Pt-Ni} , and N_{Pt-Vac} , measurements of the average numbers of Pt-Pt bonds, Pt-Ni bonds, and Pt-vacancy bonds around a Pt atom, are almost identical for the cluster expansion, parabolic potential, and constant activation energy model, with the maximum difference being about 0.19 (Table S2). These values are also consistent with our previous KMC work using the constant activation energy model [1,90]. The KRA broken bond model has the largest value of $N_{Pt-Pt} + N_{Pt-Ni}$, consistent with the less loss of octahedral shape

(Figure 5e). It is worth noting that although the parabolic potential model has a prediction error similar to that of the constant activation energy model, its Ni composition profile is much closer to that of the cluster expansion in the KMC simulations. This could be attributed to the more physically meaningful formalism of the parabolic potential model.

The simple broken bond model is the least accurate of the models, but it is much faster than the other models for calculating activation energies (by about 3 orders of magnitude) since it only evaluates number of nearest-neighbor bonds, which is equivalent to a cluster expansion with only nearest-neighbor pair ECI (Table 2). All other models calculate transition state energies roughly equally quickly, as they all use the complete cluster expansion, with 1097 ECI, to evaluate the energy difference between two different states (either the initial state and end state or the initial state and transition state).

Table 2. Average execution time for calculating the activation energy expressed relative to the time for the cluster expansion model.

	Relative time calculating activation energy
Broken bond	1.673×10^{-3}
Broken bond (KRA)	1.210
Parabolic potential	1.175
Constant activation energy	1.160
Cluster expansion	1

The overall speed of the KMC simulation is not solely determined by the time required to calculate the activation energy. In the KMC routine the total time consists of the time i) choosing a random event, ii) checking if the event is valid according to its neighboring environment, iii) calculating

the activation energy and transition probability if the event is valid, iv) updating the nanoparticle configuration if the event is accepted, and v) recording simulation data. Once calculating the activation energy becomes sufficiently fast, the remaining steps may become the bottleneck. In our implementation, due largely to the constraints on the allowed transitions on the surface, steps i) and ii) take up about as much time as calculating the activation energy using the cluster expansion, so methods that calculate the activation energy significantly more quickly than the cluster expansion have relatively little effect on the overall simulation time.

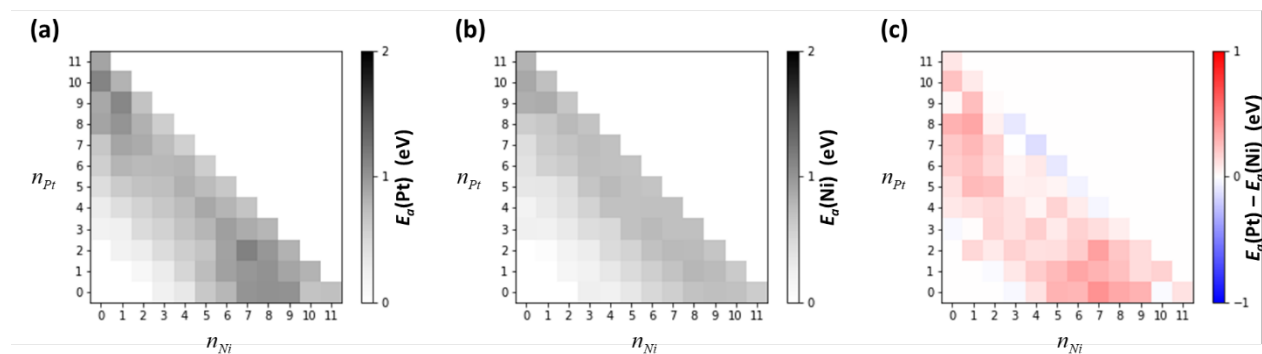


Figure 6. Cluster-expansion-predicted activation energies (E_a), averaged over 20,623 hops in randomly-generated nanoparticles, as a function of number of nearest-neighbor atoms n_{Ni} and n_{Pt} around the diffusing species. a) The activation energies for Pt. b) The activation energies for Ni. c) The differences between the activation energies in a) and b).

The transition state cluster expansion provides some insights into the activation energies for Pt and Ni diffusion as a function of the local environment. We randomly generated 100 nanoparticles with between 114 and 387 atoms and compositions ranging from 3% Ni to 97% Ni. The structures of these particles were determined in the same way as for the training structures. For each atom in each particle, we randomly generated a hop to a nearest-neighbor site (after making that site vacant if necessary), for a total of 20,623 random hops. The coordination number of the diffusing species in the initial site ranges from 3 to 11 in this data set. The average activation energies for Pt and

Ni diffusion as a function of the nearest-neighbor environment are shown in Figure 6. The largest activation energies are not observed at the largest coordination number (11) as might be expected, but rather at coordination numbers of 10 or 9. This may be a reason for the relatively poor performance of the simple broken bond model, which cannot capture this non-linear trend. The largest activation energies occur near extreme Pt/Ni (or Ni/Pt) ratios, i.e. far away from 1:1, for both Pt diffusion and Ni diffusion. This may be due to the fact that the local environment has more structural degrees of freedom when it contains a mix of Pt and Ni atoms, providing more options for finding a low-energy path. Overall, we predict that the activation energy for Pt diffusion from undercoordinated sites is larger than that for Ni, but the activation energies are closer in the bulk (Figure 6c). The relatively slow Pt diffusion near the surface may play a role in protecting the underlying Ni atoms from dissolution [1,90,94,95].

IV. CONCLUSIONS

We have presented and evaluated a cluster expansion model for predicting activation energies for vacancy-mediated diffusion in alloys by explicitly including sites representing saddle points on the potential energy surface. This approach allows for the calculation of activation energies using a single global cluster expansion and can be systematically improved with additional training data to the point at which it is roughly twice as accurate as commonly-used, simpler models, with comparable overall execution speed. Of the simpler models, the broken bond model is fastest at calculating activation energy but produces anomalous results due to its lack of detailed balance. A version of the broken bond model that corrects for this deficiency yields significantly improved results but is slower as it requires the evaluation of the energies of the initial and final states. Similar accuracy and speed are achieved by the constant activation energy model and the parabolic potential model. One benefit of the constant activation energy model is that it does not need to be

trained on DFT-calculated activation energies to predict dynamics, as it calculates relative rates only using the difference in energy between the initial and final states. However, of the simpler models, the kinetic evolution of a Pt-Ni nanoparticle is most accurately represented by the parabolic potential model. This may be due to the better physical motivation behind the parabolic potential model, but the extent to which this result depends on the particular system being studied is not clear. A weighted average of the cluster expansion model and parabolic potential model is most accurate at all training set sizes, suggesting this may be an effective approach in general.

ACKNOWLEDGEMENTS

This work was supported by the Office of Naval Research (Grant No. ONR MURI N00014-15-1-2681). C.L. and T.M. acknowledge the computational resources provided by the Brookhaven National Laboratory under grant 33818, the Maryland Advanced Research Computing Center (MARCC), and the Homewood High Performance Cluster (HHPC). The authors also thank Alberto Hernandez for helpful discussions. Atomic-scale structural images were generated using VESTA [96].

APPENDIX

In this appendix we will prove that if configuration space is constrained to prevent the simultaneous occupation of two sites, then all cluster functions that depend on the occupancies of both of those sites can be removed from the expansion and the remaining functions form a complete basis.

We start by introducing some background and notation. Let the number of possible states for the j^{th} site of a cluster expansion be given by N_j . For example, in the Pt-Ni-Vacancy cluster expansion, $N_j = 3$ for all sites, as each site can be occupied by a vacancy, Pt atom, or Ni atom. For a given configuration, let the occupation of the j^{th} site be indexed by the site variable s_j , which can take on values from 1 to N_j . For example, we may have $s_j = 1$ if the j^{th} site is vacant, $s_j = 2$ if the j^{th} site is occupied by Pt, and $s_j = 3$ if the j^{th} site is occupied by Ni. The set of site variables for all sites for a given atomic configuration is given by the vector \mathbf{s} . Here we will always define the site variables so that $s_j = 1$ whenever the j^{th} site is vacant.

To construct the cluster expansion basis, we select a set of N_j independent basis functions, $\Theta_i(s_j)$, for each site, where $i \in \{1, 2, \dots, N_j\}$. Let one of these basis functions be the constant value, 1. We will always assign this value to the first basis function, so for every site we have $\Theta_1(s_j) = 1$. The tensor product of all of these single-site bases define a complete basis of “cluster functions” that can be used to represent any function of \mathbf{s} . Every one of the cluster functions in this basis is a product of a unique combination of single-site basis functions, with the product including exactly one basis function for each site:

$$\Phi_{\mathbf{b}}(\mathbf{s}) = \prod_j \Theta_{b_j}(s_j), \quad (1)$$

where the vector \mathbf{b} represents the particular combination of single-site basis functions that define this cluster function, and b_j is the index of the single-site basis function for the j^{th} site. As $\Theta_1(s_j) = 1$, each cluster function is a function of only the site variables for which $b_j \neq 1$ [32]. The

corresponding sites are the “cluster” represented by the cluster function, and we say the cluster function “includes” these sites.

For each site, the allowed values for b_j are identical to the allowed values for s_j , as they are both integers between 1 and N_j . As the number of possible configurations is given by the number of possible unique vectors \mathbf{s} , and the number of basis functions is given by the number of possible unique vectors \mathbf{b} , we can see that the number of basis functions equals the number of total possible configurations. I.e. for any configuration \mathbf{s} , we can identify a unique basis function for which $b_j = s_j$ for all j , and vice versa.

Now we consider the case where there are two sites that cannot be occupied at the same time. We will let $j = 1$ for the first site and $j = 2$ for the second. The constraint that both sites cannot be occupied eliminates all configurations for which $s_1 \neq 1$ and $s_2 \neq 1$. As a result, there are now more cluster functions than possible configurations.

We can eliminate redundant cluster functions by eliminating all cluster functions that are functions of s_1 and s_2 , i.e. cluster functions for which $b_1 \neq 1$ and $b_2 \neq 1$. The number of remaining vectors \mathbf{s} is equal to the number of remaining vectors \mathbf{b} , as for any configuration \mathbf{s} , we can still identify a unique basis function for which $b_j = s_j$ for all j , and vice versa. Thus, the number of remaining cluster functions equals the number of allowed configurations after imposing the constraint.

To show that the remaining cluster functions form a complete basis, we need only show that they are linearly independent. We do this by demonstrating that each of the cluster functions we have removed can be expressed as a linear combination of cluster functions that were not removed. Specifically, let a cluster function that includes sites 1 and 2 be labeled by

$$\mathbf{b}_{++} = \{b_1, b_2, b_3, b_4, \dots\}, \quad (2)$$

where $b_1 \neq 1$ and $b_2 \neq 1$. We similarly define the labels

$$\begin{aligned} \mathbf{b}_{-+} &= \{1, b_2, b_3, b_4, \dots\}, \\ \mathbf{b}_{+-} &= \{b_1, 1, b_3, b_4, \dots\}, \\ \mathbf{b}_{--} &= \{1, 1, b_3, b_4, \dots\} \end{aligned} \quad (3)$$

for cluster functions in which site 1 is removed from $\Phi_{\mathbf{b}_{++}}(\mathbf{s})$, site 2 is removed from $\Phi_{\mathbf{b}_{++}}(\mathbf{s})$, and both sites 1 and 2 are removed from $\Phi_{\mathbf{b}_{++}}(\mathbf{s})$, respectively. As none of these three cluster functions include both sites 1 and 2, none will be removed from the expansion. For the allowed values of s_1 and s_2 , $\Phi_{\mathbf{b}_{++}}(\mathbf{s})$ can be written as a linear function of $\Phi_{\mathbf{b}_{-+}}(\mathbf{s})$, $\Phi_{\mathbf{b}_{+-}}(\mathbf{s})$, and $\Phi_{\mathbf{b}_{--}}(\mathbf{s})$:

$$\begin{aligned} \Phi_{\mathbf{b}_{++}}(\mathbf{s}) &= \Theta_{b_1}(s_1)\Phi_{\mathbf{b}_{-+}}(\mathbf{s}) \\ &\quad + \Theta_{b_2}(s_2)\Phi_{\mathbf{b}_{+-}}(\mathbf{s}) \\ &\quad - \Theta_{b_1}(s_1)\Theta_{b_2}(s_2)\Phi_{\mathbf{b}_{--}}(\mathbf{s}) \end{aligned} \quad (4)$$

If there are other pairs of sites that cannot be simultaneously occupied, then the procedure can be repeated.

-
- [1] L. Cao, Z. Zhao, Z. Liu, W. Gao, S. Dai, J. Gha, W. Xue, H. Sun, X. Duan, X. Pan, T. Mueller, and Y. Huang, *Matter* **1**, 1567 (2019).
- [2] D. Sun, Y. Wang, K. J. T. Livi, C. Wang, R. Luo, Z. Zhang, H. Alghamdi, C. Li, F. An, B. Gaskey, T. Mueller, and A. S. Hall, *ACS Nano* **13**, 10818 (2019).
- [3] W. Fang, J. Chen, P. Pedevilla, X.-Z. Li, J. O. Richardson, and A. Michaelides, *Nature Communications* **11**, 1689 (2020).
- [4] K. Kawahara, Y. Ohko, T. Tatsuma, and A. Fujishima, *PCCP* **5**, 4764 (2003).
- [5] L. Kunz, F. M. Kuhn, and O. Deutschmann, *The Journal of Chemical Physics* **143**, 044108 (2015).
- [6] S. P. Ong, V. L. Chevrier, G. Hautier, A. Jain, C. Moore, S. Kim, X. Ma, and G. Ceder, *Energy & Environmental Science* **4**, 3680 (2011).
- [7] C. Wang, K. Aoyagi, P. Wisesa, and T. Mueller, *Chem. Mater.* **32**, 3741 (2020).
- [8] M. S. Islam and C. A. J. Fisher, *Chem. Soc. Rev.* **43**, 185 (2014).
- [9] W. Li, Y. Yang, G. Zhang, and Y.-W. Zhang, *Nano Lett.* **15**, 1691 (2015).
- [10] R. Haugrud and T. Norby, *Nature Materials* **5**, 193 (2006).
- [11] P. Wisesa, C. Li, C. Wang, and T. Mueller, *RSC Advances* **9**, 31999 (2019).
- [12] C. Mitra, T. Meyer, H. N. Lee, and F. A. Reboredo, *The Journal of Chemical Physics* **141**, 084710 (2014).
- [13] K.-D. Kreuer, *Chem. Mater.* **8**, 610 (1996).
- [14] K. W. Edmonds, P. Bogusławski, K. Y. Wang, R. P. Campion, S. N. Novikov, N. R. S. Farley, B. L. Gallagher, C. T. Foxon, M. Sawicki, T. Dietl, M. Buongiorno Nardelli, and J. Bernholc, *Phys. Rev. Lett.* **92**, 037201 (2004).
- [15] A. A. Istratov, C. Flink, H. Hieslmair, E. R. Weber, and T. Heiser, *Phys. Rev. Lett.* **81**, 1243 (1998).
- [16] M. Ceriotti, S. Cereda, F. Montalenti, L. Miglio, and M. Bernasconi, *Physical Review B* **79**, 165437 (2009).
- [17] H. J. Jung, D. Kim, S. Kim, J. Park, V. P. Dravid, and B. Shin, *Adv. Mater.* **30**, 1802769 (2018).
- [18] A. B. Bortz, M. H. Kalos, and J. L. Lebowitz, *Journal of Computational Physics* **17**, 10 (1975).
- [19] P. D. Kolokathis and D. N. Theodorou, *The Journal of Chemical Physics* **137**, 034112 (2012).
- [20] G. H. Vineyard, *J. Phys. Chem. Solids* **3**, 121 (1957).
- [21] G. H. Gilmer, *Science* **208**, 355 (1980).
- [22] J. Erlebacher, M. J. Aziz, A. Karma, N. Dimitrov, and K. Sieradzki, *Nature* **410**, 450 (2001).
- [23] J. Erlebacher, *J. Electrochem. Soc.* **151**, C614 (2004).
- [24] B. Puchala, S.-k. Lin, L. Wang, and D. Morgan, *ECS Transactions* **50**, 1643 (2013).
- [25] Y. Tian, W. Jiao, P. Liu, S. Song, Z. Lu, A. Hirata, and M. Chen, *Nature Communications* **10**, 5249 (2019).
- [26] P. Meakin and K. M. Rosso, *The Journal of Chemical Physics* **129**, 204106 (2008).
- [27] T. H. Lim, D. McCarthy, S. C. Hendy, K. J. Stevens, S. A. Brown, and R. D. Tilley, *ACS Nano* **3**, 3809 (2009).
- [28] L. Li, P. N. Plessow, M. Rieger, S. Sauer, R. S. Sánchez-Carrera, A. Schaefer, and F. Abild-Pedersen, *J. Phys. Chem. C* **121**, 4261 (2017).
- [29] R. A. Marcus, *The Journal of Chemical Physics* **24**, 966 (1956).
- [30] P. Xiao and G. Henkelman, *ACS Nano* **12**, 844 (2018).
- [31] P. Lindgren, G. Kastlunger, and A. A. Peterson, *arXiv: 1903.09903 [physics. chem-ph]* (2019).
- [32] J. M. Sanchez, F. Ducastelle, and D. Gratias, *Physica A: Statistical Mechanics and its Applications* **128**, 334 (1984).
- [33] A. Van der Ven, G. Ceder, M. Asta, and P. D. Tepesch, *Physical Review B* **64**, 184307 (2001).

- [34] S. Sawarkar and A. Chatterjee, Computational Materials Science **184**, 109876 (2020).
- [35] H.-J. Lu, N. Zou, R. Jacobs, B. Afflerbach, X.-G. Lu, and D. Morgan, Computational Materials Science **169**, 109075 (2019).
- [36] L. Messina, A. Quaglino, A. Goryaeva, M.-C. Marinica, C. Domain, N. Castin, G. Bonny, and R. Krause, arXiv: 1808.06935 [physics. chem-ph] (2018).
- [37] H. Wu, A. Lorensen, B. Anderson, L. Witteman, H. Wu, B. Meredig, and D. Morgan, Computational Materials Science **134**, 160 (2017).
- [38] W. Li, Y. Ando, E. Minamitani, and S. Watanabe, The Journal of Chemical Physics **147**, 214106 (2017).
- [39] S. Kondati Natarajan and J. Behler, J. Phys. Chem. C **121**, 4368 (2017).
- [40] X.-S. Kong, K.-N. He, J. Hou, T. Zhang, and C. S. Liu, <http://dx.doi.org/10.2139/ssrn.3406950> (2019).
- [41] Y. Zeng, Q. Li, and K. Bai, Computational Materials Science **144**, 232 (2018).
- [42] A. Hernandez, A. Balasubramanian, F. Yuan, S. A. M. Mason, and T. Mueller, npj Computational Materials **5**, 112 (2019).
- [43] C. M. Van Baal, Physica **64**, 571 (1973).
- [44] Z. W. Lu, B. M. Klein, and A. Zunger, Modell. Simul. Mater. Sci. Eng. **3**, 753 (1995).
- [45] A. Van der Ven, M. K. Aydinol, G. Ceder, G. Kresse, and J. Hafner, Physical Review B **58**, 2975 (1998).
- [46] B. Sadigh, M. Asta, V. Ozoliņš, A. K. Schmid, N. C. Bartelt, A. A. Quong, and R. Q. Hwang, Phys. Rev. Lett. **83**, 1379 (1999).
- [47] A. van de Walle and G. Ceder, Journal of Phase Equilibria **23**, 348 (2002).
- [48] O. Wieckhorst, S. Müller, L. Hammer, and K. Heinz, Phys. Rev. Lett. **92**, 195503 (2004).
- [49] B. C. Han, A. Van der Ven, G. Ceder, and B.-J. Hwang, Physical Review B **72**, 205409 (2005).
- [50] K. Yuge, J. Phys.: Condens. Matter **22**, 245401 (2010).
- [51] T. Mueller and G. Ceder, ACS Nano **4**, 5647 (2010).
- [52] R. V. Chepulskii, W. H. Butler, A. van de Walle, and S. Curtarolo, Scripta Mater. **62**, 179 (2010).
- [53] X. Huang, Z. Zhao, L. Cao, Y. Chen, E. Zhu, Z. Lin, M. Li, A. Yan, A. Zettl, Y. M. Wang, X. Duan, T. Mueller, and Y. Huang, Science **348**, 1230 (2015).
- [54] T. Mueller, Physical Review B **86**, 144201 (2012).
- [55] C. Li, D. Raciti, T. Pu, L. Cao, C. He, C. Wang, and T. Mueller, J. Phys. Chem. C **122**, 18040 (2018).
- [56] L. Cao and T. Mueller, Nano Lett. **16**, 7748 (2016).
- [57] T. L. Tan, L.-L. Wang, D. D. Johnson, and K. Bai, Nano Letters **12**, 4875 (2012).
- [58] M. K. Y. Chan, J. Reed, D. Donadio, T. Mueller, Y. S. Meng, G. Galli, and G. Ceder, Physical Review B **81**, 174303 (2010).
- [59] A. Van der Ven and G. Ceder, Phys. Rev. Lett. **94**, 045901 (2005).
- [60] T. Rehman, M. Jaipal, and A. Chatterjee, Journal of Computational Physics **243**, 244 (2013).
- [61] F. Soisson and C.-C. Fu, Physical Review B **76**, 214102 (2007).
- [62] F. Soisson and C. C. Fu, Solid State Phenomena **129**, 31 (2007).
- [63] Y. L. Bouar and F. Soisson, Physical Review B **65**, 094103 (2002).
- [64] A. Van der Ven, J. C. Thomas, Q. Xu, B. Swoboda, and D. Morgan, Physical Review B **78**, 104306 (2008).
- [65] J. Bhattacharya and A. Van der Ven, Physical Review B **81**, 104304 (2010).
- [66] G. Henkelman, B. P. Uberuaga, and H. Jónsson, The Journal of Chemical Physics **113**, 9901 (2000).
- [67] G. Henkelman and H. Jónsson, The Journal of Chemical Physics **113**, 9978 (2000).
- [68] V. R. Stamenkovic, B. Fowler, B. S. Mun, G. Wang, P. N. Ross, C. A. Lucas, and N. M. Marković, Science **315**, 493 (2007).

- [69] L. Cao, L. Niu, and T. Mueller, *Proceedings of the National Academy of Sciences* **116**, 22044 (2019).
- [70] C. Cui, L. Gan, M. Heggen, S. Rudi, and P. Strasser, *Nature Materials* **12**, 765 (2013).
- [71] X. Huang, E. Zhu, Y. Chen, Y. Li, C.-Y. Chiu, Y. Xu, Z. Lin, X. Duan, and Y. Huang, *Adv. Mater.* **25**, 2974 (2013).
- [72] W. Kohn and L. J. Sham, *Physical Review* **140**, A1133 (1965).
- [73] G. Kresse and J. Furthmüller, *Physical Review B* **54**, 11169 (1996).
- [74] J. P. Perdew, K. Burke, and M. Ernzerhof, *Phys. Rev. Lett.* **77**, 3865 (1996).
- [75] B. Hammer, L. B. Hansen, and J. K. Nørskov, *Physical Review B* **59**, 7413 (1999).
- [76] P. E. Blöchl, *Physical Review B* **50**, 17953 (1994).
- [77] P. Wisesa, K. A. McGill, and T. Mueller, *Physical Review B* **93**, 155109 (2016).
- [78] M. Methfessel and A. T. Paxton, *Physical Review B* **40**, 3616 (1989).
- [79] E. Ising, *Zeitschrift für Physik* **31**, 253 (1925).
- [80] D. B. Laks, L. G. Ferreira, S. Froyen, and A. Zunger, *Physical Review B* **46**, 12587 (1992).
- [81] L. Cao, C. Li, and T. Mueller, *Journal of Chemical Information and Modeling* **58**, 2401 (2018).
- [82] Q. Wu, B. He, T. Song, J. Gao, and S. Shi, *Computational Materials Science* **125**, 243 (2016).
- [83] T. Mueller and G. Ceder, *Physical Review B* **80**, 024103 (2009).
- [84] J. S. Lim, N. Molinari, K. Duanmu, P. Sautet, and B. Kozinsky, *J. Phys. Chem. C* **123**, 16332 (2019).
- [85] Y. Tiwary and K. A. Fichthorn, *Physical Review B* **81**, 195421 (2010).
- [86] F. Montalenti and R. Ferrando, *Phys. Rev. Lett.* **82**, 1498 (1999).
- [87] See Supplemental Material at [URL will be inserted by publisher] for additional details about constructing and fitting of the cluster expansion and details about the Monte Carlo simulations.
- [88] N. Metropolis, A. W. Rosenbluth, M. N. Rosenbluth, A. H. Teller, and E. Teller, *The Journal of Chemical Physics* **21**, 1087 (1953).
- [89] L. Cao and T. Mueller, *J. Phys. Chem. C* **119**, 17735 (2015).
- [90] Q. Jia, Z. Zhao, L. Cao, J. Li, S. Ghoshal, V. Davies, E. Stavitski, K. Attenkofer, Z. Liu, M. Li, X. Duan, S. Mukerjee, T. Mueller, and Y. Huang, *Nano Lett.* **18**, 798 (2018).
- [91] D. H. Wolpert, *Neural Networks* **5**, 241 (1992).
- [92] J. Hartung, G. Knapp, and B. K. Sinha, *Statistical meta-analysis with applications* (John Wiley & Sons, 2008).
- [93] C. Cui, L. Gan, M. Heggen, S. Rudi, and P. Strasser, *Nature Materials* **12**, 765 (2013).
- [94] V. Beermann, M. Gocyla, S. Kühn, E. Padgett, H. Schmies, M. Goerlin, N. Erini, M. Shviro, M. Heggen, R. E. Dunin-Borkowski, D. A. Muller, and P. Strasser, *J. Am. Chem. Soc.* **139**, 16536 (2017).
- [95] V. Beermann, M. Gocyla, E. Willinger, S. Rudi, M. Heggen, R. E. Dunin-Borkowski, M.-G. Willinger, and P. Strasser, *Nano Lett.* **16**, 1719 (2016).
- [96] K. Momma and F. Izumi, *J. Appl. Crystallogr.* **41**, 653 (2008).