

Active learning and molecular dynamics simulations to find high melting temperature alloys

David E. Farache^{a,*}, Juan C. Verduzco^{a,*}, Zachary D. McClure^a,
Saaketh Desai^b, and Alejandro Strachan^a

^a School of Materials Engineering and Birck Nanotechnology Center
Purdue University, West Lafayette, Indiana, 47907 USA

^b Center for Integrated Nanotechnologies
Sandia National Laboratories, Albuquerque, New Mexico, 87123 USA
* Authors contributed equally to this work.

Abstract

Active learning (AL) can drastically accelerate materials discovery; its power has been shown in various classes of materials and target properties. Prior efforts have used machine learning models for the optimal selection of physical experiments or physics-based simulations. However, the latter efforts have been mostly limited to the use of electronic structure calculations and properties that can be obtained at the unit cell level and with negligible noise. We couple AL with molecular dynamics simulations to identify multiple principal component alloys (MPCAs) with high melting temperatures. Building on cloud computing services through nanoHUB, we present a fully autonomous workflow for the efficient exploration of the high dimensional compositional space of MPCAs. We characterize how uncertainties arising from the stochastic nature of the simulations and the acquisition functions used to select simulations affect the convergence of the approach. Interestingly, we find that relatively short simulations with significant uncertainties can be used to efficiently find the desired alloys as the random forest models used for AL average out fluctuations.

Keywords - multiple principal component alloys (MPCAs); active learning; uncertainty quantification;

1 Introduction

The combination of experiments, physics-based simulations, and data science tools has been shown to have the potential to accelerate discovery of novel materials with optimized properties.¹⁻¹² Active learning (AL), a subset of machine learning in which models learn dynamically, has gathered significant interest, both from the point of basic science¹³ and for commercial applications.¹⁴ AL models analyze existing data and formulate queries to acquire additional information towards a design goal. AL workflows start with a model trained with an initial set of data and evaluate the expected gain towards an objective function of all possible new experiments within a design space.¹⁵ Top candidates are characterized (e.g. by performing an experiment) and the outcome is added to

the existing dataset. Given this additional information, a new cycle is started. With this iterative process, illustrated in Figure 1, the model becomes more accurate in regions of interest within the design space. In order to identify the next best query, AL uses selection strategies known as information acquisition functions. These strategies differ from each other in the relative balance between exploitation and exploration included in their mathematical formulation. Exploitation favors cases expected to maximize the objective function while exploration focuses on areas of high uncertainty.




Figure 1: Schematic representation of the AL iterative optimization. Starting from existing data, a predictive model is trained with machine learning. The model is then applied to all candidate materials in the design space to assess their performance, including uncertainties. An acquisition function is used to select the next material(s) to be tested (either via experiments or simulations). If the new material does not satisfy the design conditions, it is added back to the existing data set and the cycle re-initiated.

AL has been used for a wide range of applications from natural language processing,¹⁶ reaction screening for pharmaceutical applications,¹ and multiscale modeling.^{2,17} In materials science, AL has been used to accelerate discovery of materials with desired properties by coupling it with experiments^{4,18–22} and physics-based simulations.^{5–7} In addition, AL workflows paired with existing closed data sets has shown the ability of these models to reduce the number of queries needed to identify the best candidate.^{9–12}

Kusne et al.³ developed a closed-loop autonomous algorithm for materials discovery and used it to explore the Ge-Sb-Te ternary system in search for an optimal phase-change material with the highest optical bandgap difference between phases. With an iterative approach using experimental ellipsometry data, they identified a new composition, $\text{Ge}_4\text{Sb}_6\text{Te}_7$, with nearly three times the optical bandgap of $\text{Ge}_2\text{Sb}_2\text{Te}_5$, a material used in random-access memory devices. Xue et al.⁴ demonstrated an adaptive design strategy for the search of NiTi-based shape memory alloys with low thermal hysteresis. Starting with 26 materials synthesized their optimization algorithm was used to screen a space of $\sim 800,000$ candidate materials. This screening resulted in 36 alloys queried and tested

using differential scanning calorimetry (DSC) with several showing lower thermal hysteresis than the starting training set. For additional examples, see Refs.^{19–22}

Coupling AL workflows with physics-based simulations can further reduce the costs and time required to discover new materials by minimizing the number of experimental trials required. Tran and Ulissi⁵ explored a high dimensional space of density functional theory (DFT) results on inter-metallics in search for optimized catalyst materials for CO₂ reduction. Their approach produced a set of 54 promising materials from ~1499 candidates in the Materials Project.⁶ Out of these materials, some have been since confirmed experimentally. Seko et al.⁷ reported a virtual screening of 54,779 materials looking for compounds of low lattice thermal conductivity. They used AL to select appropriate descriptors for a Gaussian process regression (GPR) model used to get information on the large dataset. Their screening yielded 221 materials with very low conductivity. After some constraints were imposed, 19 of them were characterized using DFT calculations for possible applications as thermoelectric materials.

Work involving the combination of AL or similar optimization methods and materials simulations has been mostly limited to DFT calculations and to properties that can be obtained from relatively small simulation cells without noise. While DFT often offers a good balance between accuracy and computational efficiency, it remains limited to a subset of materials properties. Fortunately, materials models across scales are available²³ and significant progress in their coupling across scales has occurred over the last few decades. Examples range from crystal plasticity models^{10,24} to interatomic potentials from ab-initio simulations.^{17,25–29} Physics-based materials models across scales open the possibility of significantly expanding the reach of AL approaches.

Molecular dynamics (MD) is an important rung in multiscale modeling that connects *ab initio* physics to the meso- and macro-scales. MD has been paired with optimization methods like reinforcement learning, including Monte Carlo tree search (MCTS).^{30,31} For example, Patra et al.³² paired MCTS with MD to find a copolymer compatibilizer, and Loeffler et al.³³ applied the methods to study defect structures in metal dichalcogenides.

In this work, we couple AL with MD simulations to find multiple principal component alloys (MPCAs) with high predicted melting temperature for a model CrCoCuFeNi system. The main objective of this paper is to demonstrate the potential of coupling of AL workflows with MD simulations and characterizing the effect of the stochastic and noisy nature of MD results on AL workflows. We find that even under the noise in relatively short MD simulations, acquisition functions that balance exploration with exploitation are capable of finding the desired alloy efficiently. This paper is organized as follows: Section 2 introduces the overall approach, design space, and simulation details. Section 3 describes the AL runs exploring the effects of (a) different information acquisition functions and (b) different MD simulation times. It also discusses how uncertainty of the model evolves as the active learning workflow progresses. In Section 4 we summarize and draw final conclusions of the paper. Finally, in Section 5 we discuss availability of our models, datasets and workflows through nanoHUB.

2 Methods

2.1 Problem statement

MPCAs are a new class of materials, which includes high entropy and complex concentrated alloys, where four or more elements are combined in nearly equal atomic percentages. Interest in these materials has grown due to several desirable properties. Properties such as high strength at elevated temperatures, radiation resistance,³⁴ and high melting temperatures make some MPCAs attractive

for high temperature and extreme applications.^{35–37} Experimental determination of this melting temperature is challenging because of their complex processing³⁸ protocols, chemical reactions, and phase separation.^{39,40} For these reasons, using computational methods to determine the melting temperature of these alloys is highly desirable. Fortunately, melting temperature can be accurately obtained from first principles, and the first calculations for MPCAs are emerging.⁴¹

The high-dimensional space of possible compositions prevents brute-force approaches even for relatively cheap simulations. Therefore, an efficient way to explore this space is essential. In this work, we pair AL with MD simulations to find MPCAs with the highest possible temperature in a model system. AL in the context of materials design seeks to reduce the number of experiments needed to find an optimal candidate out of a pool of untested contenders in an unexplored space. Within the AL scheme, predictions of the surrogate machine learning model, along with sample-wise uncertainties, are used to identify the next candidate to be tested, via acquisition functions. In this work, we make use of the FUELS (Forests with Uncertainty Estimates for Learning Sequentially) framework proposed by Ling et al.¹¹ based on random forests and a supervised machine learning algorithm.

2.1.1 Initial and candidate space

We explore 5-component FCC alloys incorporating Cr, Co, Cu, Fe, and Ni. Iterative AL models require an initial subset of entries to be evaluated to serve as the prior for the ML model. Our initial set of alloys was selected to be around the equiatomic composition $\text{Cr}_{20}\text{Co}_{20}\text{Cu}_{20}\text{Fe}_{20}\text{Ni}_{20}$. Alloys included in this initial set deviate by 10 at.% from this composition, such that each element composes 10 at.% to 30 at.% of the alloy. Using steps of 10 at. %, we get 39 data points. Melting temperatures were calculated via MD simulations, as described below, using the two-phase coexistence method. This initial set was the starting point for all of our AL runs.

For our exploration space, we relaxed the constraints from the initial set to allow more deviations from the equiatomic composition such that any of the elements is between 0 at.% and 50 at.%. This allows for exploration of binaries, ternaries and quaternaries. Using composition steps of 10 at. %, this results in 554 possible alloys that are not included in our initial subset. This grid spacing may seem coarse, but it is appropriate for an initial exploration and could be refined in a subsequent step if needed. This refinement was not necessary as will be shown below. Our iterative AL model starts with a training step on our initial space, and the candidate space will serve as the pool from which it will be drawing candidates.

2.2 MD simulations of melting

To predict the melting temperature of FCC MPCAs, we made use of MD simulations and the two-phase coexistence method.^{42,43} To describe atomic interactions, we used a many-body, embedded-atom-method interatomic potential developed by Farkas and Caro.⁴⁴ The potential was designed for FCC Cr-Co-Cu-Fe-Ni alloys with near equiatomic compositions. Its parameterization focused on reproduction of elastic constants, vacancy formation energy, stacking fault energy, surface energies and relative phase stability to ensure that FCC is the stable structure for all five components. However, this potential has not been trained or tested for melting temperature calculations.

Several techniques exist to compute the melting temperature of materials using MD simulations. The most direct approach involves heating a crystalline sample until it melts and then cooling the melt until it recrystallizes. This results in overheating and undercooling, requiring a-posteriori corrections.⁴⁵ To avoid these issues, we predict melting temperatures using a phase coexistence

method.^{42,43} In this method, a temperature at which the crystal and liquid phases of the material of interest coexist is established.

The first step is to generate an initial sample with a liquid and solid phase in contact, we achieve this with the use of two thermostats during the sample preparation step. We start from the four-atom FCC unit cell with a lattice parameter of 3.56 Å and replicate it 8 x 8 x 18 to create a periodic simulation cell. The third direction (z) is normal to the solid-liquid interface. Atom types are assigned based on the desired composition randomly to each lattice site. We note that this approach ignores possible short-range order in these alloys. This is an acceptable compromise since our goal is to demonstrate the coupling of active learning with MD simulations and, thus, the accuracy of the model predictions is of secondary importance. This random alloy assumption was also made by the original authors of the interatomic potential⁴⁴ and, in a subsequent paper, Farkas and Caro showed that the short-range ordering effects were negligible with a similarly parametrized EAM potential for the Cr-Co-Al-Fe-Ni system.⁴⁶ The entire structure is relaxed via energy minimization with respect to cell parameters and atomic positions. To create initial liquid/solid structure the cell is then divided in halves along z and two thermostats are used to create the two phases. The two regions are equilibrated at their initial temperatures for 10 ps using isothermal, isobaric MD simulations (NPT ensemble) with a Berendsen thermostat⁴⁷ and Nose-Hoover barostat.⁴⁸ This first step seeks to create an initial cell with both liquid and solid and the two temperatures need to be chosen appropriately (the temperature of the liquid half should be higher than the melting temperature). A snapshot of the system after this initial step is shown in Figure 2. Clearly, this initial sample is not in equilibrium, and we follow the initial step with an isoenthalpic, isobaric (NPH) simulation where the liquid and solid regions are allowed to exchange energy freely and come to equilibrium. After steady state is achieved and the simulation cell has a uniform temperature, if the sample contains both liquid and solid the simulation temperature corresponds to the melting temperature of the system. The use of the NPH ensemble allows the temperature of the system to evolve towards the melting temperature. If the sample is initially too hot the liquid phase will grow at the expense of the solid, the heat of fusion will automatically cool the sample down towards the melting temperature; the reverse happens if the sample is initially too cold.




Figure 2: MD snapshot of the simulation cell divided into a liquid and solid region for an equiatomic alloy of Cr, Co, Cu, Fe and Ni. Each color represents a different element.

A successful run requires the coexistence of liquid and solid after steady state has been achieved. These conditions are checked by analyzing the atomic structure and the time evolution of the instantaneous temperature of the system. Phase fractions of solid and liquid are calculated using the polyhedral template matching (PTM) algorithm from OVITO.^{49,50} Coexistence is defined as the fraction of liquid being between 35 and 65 at.% and the FCC crystal also between 35 and 65 at.% in the final snapshot of the simulation. We also analyze thermodynamic data from the last 60% of the simulation to assess steady state, we require the absolute value of the slope of the

instantaneous temperature vs. time to be less than 1 K/ps. Finally, we also compute the 95% confidence interval on the temperature calculation. Melting temperatures for our initial set ran with 100ps of simulation time.

2.3 Uncertainties and noise in the MD predictions

The description of interatomic forces is at the heart of MD simulations and determines the predicted melting temperatures. The potential used in this work was not developed or widely tested for melting temperature calculations and we are unaware of any potential that has. This limits the accuracy of our melting temperature predictions. However, as stated above, our focus is on understanding the use of MD in ML workflows and not in accurate predictions of melting temperatures. Our work can be easily extended to MD simulations using higher-accuracy interatomic potentials^{51–53} or density functional theory calculations.⁵⁴ Nevertheless, for completeness, Figure 3 compares the predicted melting temperatures with experimental values⁵⁵ for various alloys and single-elements.⁵⁶ Our simulations assume random FCC alloys since that is what the AL workflow uses, this does not match the experiments in all cases. We included experimental melting temperatures for single elements for comparison. All MD simulations all arrange structures as FCC. For iron, we show the predicted melting temperature of FCC Fe since the value is relevant for the high-melting temperature alloys found by the AL workflow. Interestingly, for Cobalt (Co), the MD simulation with this potential achieved coexistence by creating a solid BCC structure rather than FCC. We find that there is a consistent overestimation of melting temperatures, except for Cu and Cr.




Figure 3: Comparison for MD simulated and experimental melting temperatures for alloys composed of elements included in the potential. Dashed line indicates a match between the values. Experimental values for MPCAs taken from Wu et al.⁵⁵ Experimental values for single-elements taken from NIST.⁵⁶ Filled symbols represent alloys reported as FCC crystal structure, open symbols represent non-FCC crystal structures.

It is important to note that the generation of the initial atomic structure is done stochastically, and temperature is obtained as the time-average of a fluctuating quantity. Thus, the resulting

melting temperature from independent runs will vary due to sample-to-sample variability (which can be reduced increasing the simulation cell size) and time averaging of instantaneous temperature (which can be reduced by running longer simulations and also by increasing system size). One of the goals of this paper is to assess what level of noise can be tolerated in an AL workflow.

To quantify the variability in predicted melting temperatures, we chose one of our optimal alloy candidates. We performed 36 simulations with the same overall composition but independent atomic configurations and initial velocities. Figure 4 shows the resulting distributions of predicted melting temperatures for the three simulation times used in this study, the inset shows the mean and uncertainty estimates. Figure 4 shows that, as expected, repeated experiments result in mean values relatively independent of the simulation time (2470-2480 K) and reducing the simulation time results in broader distributions (remember that temperature is associated with the time-averaged kinetic energy per degree of freedom).




Figure 4: Distributions of MD simulated melting temperatures for a MPCA of 50% Fe and 50% Co at different simulation times. Distributions indicate a mean temperature around 2473 K and narrow down as the simulation time increases.

Table 1: Mean and uncertainty estimates for the distributions in Figure 4 of MD simulated melting temperatures for a MPCA of 50% Fe and 50% Co at different simulation times. These distributions were constructed by running 36 simulations with combinations of 6 different atomic arrangements and 6 different atomic velocities initializations.

TIME PER SIMULATION (PS)	MEAN MD SIMULATED TM (K)	STANDARD DEVIATION MD SIMULATED TM (K)
50	2479	17
100	2476	13
200	2470	8

2.4 Random forests and uncertainties

Several regression models, including neural networks, gaussian processes, and random forests have been used in AL workflows. Neural networks have stark advantages in image recognition and other high-dimensionality problems and have seen widespread adoption in a myriad of problems and disciplines. However, uncertainty quantification on their applications for regression and classification is remains an active area of research.⁵⁷ Random forests, on the other hand, are often preferred when dealing with tabulated data with well-defined and limited descriptors.⁵⁸ Random forests also have the ability to yield descriptor importance for interpretability⁵⁹ and there has been significant work on quantification of uncertainties.^{11,60,61}

In this work, we selected random forests as uncertainties are critical to make use of information acquisition functions within AL and to deal with noisy data present within our simulations. We use the lolopy library^{62,63} to develop RF models for melting temperature of MPCAs using composition and derived descriptors as inputs. The models provide an expectation value and sample-wise uncertainties that are fed to various acquisition functions to select the next experiment. Uncertainty estimates are obtained based on the recalibrated bootstrap forest variance as formulated in lolopy version 1.1.1.⁶³

2.5 Materials Descriptors and Model Hyperparameters

For the RF models described in subsection 2.4 to predict melting temperature, it is vital to represent our alloys with descriptors that can serve as inputs. A brute-force approach would use composition as the sole inputs but, given the relative scarcity of the data in materials applications, enhancing the inputs with descriptors or features expected to correlate with the outputs is highly beneficial.⁶⁴ Previous work by Zhang et al.⁶⁵ in which they used a classification algorithm to predict the phases of alloys outlined a materials descriptor space and a genetic algorithm strategy to down select a subset of the pool of properties. The complete set included 70 properties based on the molar average value and mismatch value for each elemental property of the alloy. Their work showed that RFs outperform other models in the prediction task, but their algorithm selects subsets limited to four descriptors. Similarly, work by Zhou et al.⁶⁶ shows that machine learning models can help the exploration of phase design of high entropy alloys. Their model uses 13 descriptors based on the average and standard deviation of properties like atomic radius, melting temperature, mixing enthalpy, mixing entropy, electronegativity, bulk modulus and valence electron concentration.

In this work, atomic properties for elements contained in the alloys are queried from properties available online on Pymatgen⁶⁷ and used to generate a unique fingerprint of each alloy using a weight-average rule. The final set of descriptors was based on the strongest correlation with melting temperature of our initial set, measured with Pearson correlation coefficients. This can be seen in Supplemental Information (SI) as Figure S1. However, it nicely reflects some of the candidates proposed by Zhang et al.⁶⁵ as it includes properties such as: melting point, electronegativity, boiling temperature, atomic mass, atomic radii, and density. Another subset of the highly correlated descriptors relates to mechanical and thermal properties and includes Young’s modulus, Poisson’s ratio, hardness, and coefficient of thermal expansion. Finally, electrical resistivity was added to the descriptors as it showed a high correlation to melting temperature.

To test our initial model, we split the initial set in an 80/20 train/test split. In Figure 5 we present a parity plot of the melting temperatures predicted by our RF model and the melting temperatures obtained via MD simulations.




Figure 5: Comparison for MD simulated and RF predicted melting temperatures for the 39 MPCAs compositions in our initial set.

Given the small initial dataset, large deviations in MAE can occur due to the random selection of the testing set. To properly determine the model's accuracy, the dataset was shuffled, split, and rerun 30 times through the random forest. The average of the MAE for the RF predicted melting temperature with respect to the MD simulated temperature was 36.44 K with an uncertainty estimate of 12.2 K. A 10-fold cross validation analysis was used to create a normalized residual distribution to determine calibration of the uncertainty. Models with well calibrated uncertainties would produce a Gaussian distribution with a mean of zero and unit standard deviation as seen in Figure 6.




Figure 6: *Probability densities of normalized residuals of RF model computed through tenfold cross-validation. Solid line is perfectly calibrated uncertainties.*

We optimized the number of estimators (trees) to 350 trees at the maximum depth of each estimator, using the mean absolute error (MAE) as a metric to measure the accuracy of our model predictions. Increasing the number of trees reduces variance, while increasing the maximum depth can help reduce bias. Our optimization stopped based on the idea that the larger quantity of either would improve the RF until we observed diminishing returns. Plots for the optimization of these parameters are included in the supplemental information as Figure S2.

2.6 Acquisition functions

Acquisition functions are mathematical expressions designed to select which of the possible tests (MD simulations in our case) to carry out next. This is done by considering the model predicted value and uncertainty estimate of the quantity of interest for all possible tests. For our case, we use the predicted value and uncertainty estimates provided by the RF model. Various acquisition functions have been proposed with different balances between exploitation and exploration. Exploration functions select areas of high uncertainty and, on the opposite end, functions based purely on exploitation select candidates with the highest expected values.

A purely exploitative approach would be maximizing the mean predicted value; this maximum mean (MM) function can be written as:

$$MM : x^* = \operatorname{argmax} E[M(x_i)] \quad (1)$$

Such functions can easily get trapped in local maxima, and some degree of exploration is often desirable. Upper confidence bound (UCB) queries the sample with the maximum predicted value plus its uncertainty estimate. For this study, the adjustable parameter K is set to unity ($K = 1$).

$$UCB : x^* = \operatorname{argmax} (E[M(x_i)] + K * \sigma[M(x_i)]) \quad (2)$$

Maximum likelihood of improvement (MLI) chooses the sample with the highest probability of surpassing the current best previously evaluated material. This probability can be computed as the Z-score for predictions that are normally distributed. Therefore, it is represented by the difference between the expected value of the prediction and the value of the current best case, x_{best} , over the uncertainty estimate.

$$MLI : x^* = \operatorname{argmax} \frac{E[M(x_i)] - E[M(x_{best})]}{\sigma[M(x_i)]} \quad (3)$$

Maximum expected improvement (MEI)⁶⁸ works by modeling our knowledge of the prediction as a normal distribution. It uses the model’s mean prediction and uncertainty estimate to draw a normal probability density function at each point. The improvement can then be understood as the probability that the function at this value surpasses our current maximum. Intuitively, this relates to the tail of the distribution crossing our current maximum. In the absence of uncertainties, the function maximum expected improvement (MEI),^{68,69} reduces to maximizing the expected values. We note that the MM function was incorrectly labeled as MEI in Refs.^{11,12}

$$MEI : x^* = \operatorname{argmax} \rho (E[M(x_i)] - E[M(x_{best})], \sigma[M(x_i)])$$

$$\text{with } \rho(z, s) = \begin{cases} s\phi'(\frac{z}{s}) + z\phi(\frac{z}{s}) & s > 0 \\ \max(z, 0) & s = 0 \end{cases} \quad (4)$$

Finally, maximum uncertainty (MU) is at the extreme end of exploration and only focuses on the candidates with the highest uncertainty in their predictions. MU has little incentive to find the top performer and finds its best use when trying to fine-tune ML models for surrogate-based optimization.

$$MU : x^* = \operatorname{argmax} \sigma[M(x_i)] \quad (5)$$

In these Equations, x^* marks the composition selected by the function to be tested next. x_i are the possible experiments to run (our search space). x_{best} is the current best performer in the training set. $E[M(x_i)]$ is the expectation value of the prediction at point x . This expectation value is equal to $\frac{1}{N} \sum_j^{n_T} t_j(x)$ where $t_j(x)$ is the prediction of tree j for point x and n_T is the total number of trees. Within MEI, ϕ represents the Gaussian cumulative distribution function, and ϕ' is the Gaussian probability density function. Finally, “arg max” (arguments of the maxima) operation returns the sample material where the function is maximized for the quantity of interest.

3 Active learning: exploring high melting temperature alloys

As described above, we started with knowledge of the melting temperature of 39 alloys and use AL to search for the alloy with the highest melting temperature in a 5-dimensional compositional space with 554 possible alloys. The performance of these acquisition functions was assessed from AL workflows with a budget of 40 experiments, mimicking time and resource constraints in real

world applications. This budget is somewhat arbitrary and chosen based on prior active learning efforts. Here an experiment is equal to the convergence of a composition of an alloy which may require several simulations. As previously described, convergence is defined when the simulation outputs an alloy that reached temperature equilibrium and the fraction of atoms in each of the phases needs to account for 35% to 65% of the system. At each step of the iterative process, we use the RF models to evaluate each acquisition function over all unexplored alloys and select the alloy that maximizes the selected function. A new MD simulation is launched with that alloy and the resulting data is added to the data set and the processes is restarted.

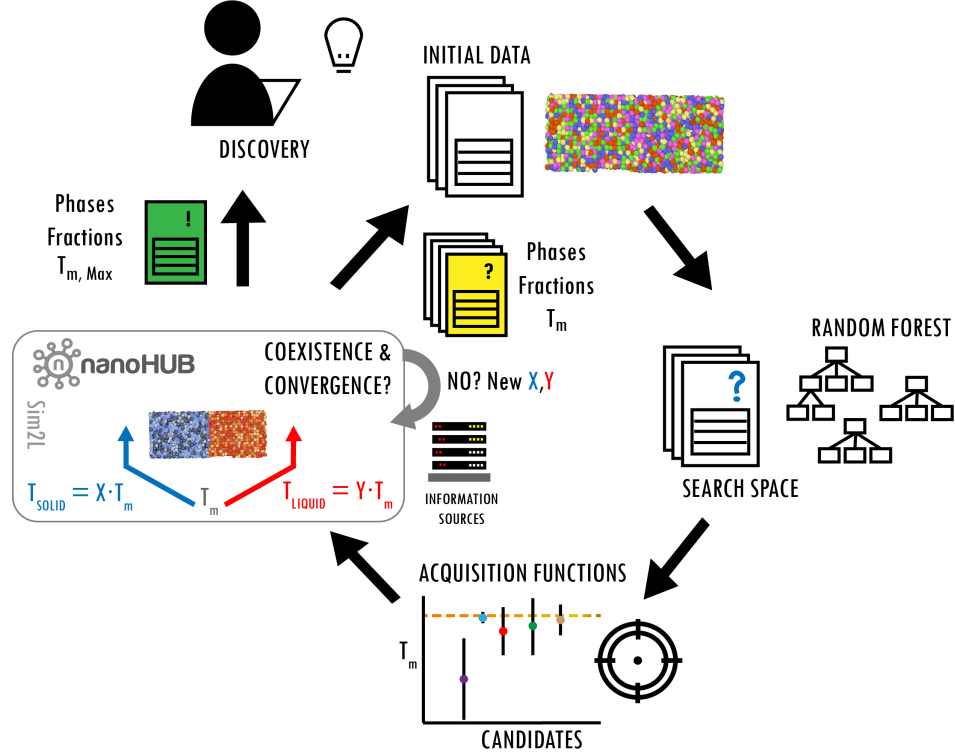


Figure 7: Schematic representation of the active learning iterative process outlined in this work. Initial data for random forest training is generated from MD simulations in a selected subspace. After training, an acquisition function is selected to provide a candidate composition for testing. Within the Simtool a new MD simulation is performed, characterized, and a secondary loop is performed to modify inputs to ensure final system convergence. Upon convergence, the data is collected, augmented to the training set, and the cycle continues until the budget is exhausted.

The MD simulation is run for the desired time and convergence, or lack thereof, is determined with the criteria in Section 2.2. Simulation inputs are box length $n=18$, MD simulation time, and temperatures of the liquid and solid regions. Additional inputs include seeds for pseudo-random number generation for initialization of the initial random structure and atomic velocities. The initial temperatures for the liquid and solid regions from the melting temperature predicted by the RF (T_m) as follows: $T_{\text{liquid},t=0} = 1.25 T_m$; $T_{\text{solid},t=0} = 0.5 T_m$. If a converged melting temperature calculation is not reached, the initial temperatures are adjusted, and a new simulation launched. The initial temperatures are increased or decreased by 5% if the resulting structural analysis indicated little liquid or solid left, respectively. This process continues until convergence is achieved. Once convergence is reached the value will be added into the known set, the model will make new predictions, and the process will repeat until our design goal is reached or our experiment

budget is exhausted.

3.1 Active Learning results

To explore the combination of AL with MD simulations we tested six selection strategies (five information acquisition functions described in Section 2.6 and a random sampling baseline). To assess the effect the noise in the data we considered three simulation times 50, 100 and 200 ps. Each of the 18 runs used the same initial set as described in section 2.1.1.

Figures 8, 9 and 10 show the melting temperature, both predicted by the RF (filled symbols) and the one obtained from the subsequent MD simulation (open symbols), as a function of experiment number for the various acquisition functions and simulation times. The insets in Figures 8, 9 and 10, show the composition of the alloys with the highest predicted melting temperatures. Tables 2, 3 and 4 summarize the various AL outcomes: for the alloy with the highest melting temperature within the budget of 40 possible compositions for each case, we list the temperature predicted by the RF and the actual value obtained from the MD simulations, the uncertainty in the RF model, number of experiments required to locate the alloy with the highest melting temperature and its composition.




Figure 8: Performance of different acquisition functions in a 40-experiment budget AL run with 50 picoseconds MD simulation time. All functions start from identical initial sets. Open symbols represent MD simulated melting temperatures. The filled symbols with error bars represent RF predicted melting temperatures. Xx represents other elements, Cu and Cr. The insert includes a close-up on the best performing MPCAs compositions. These compositions contain high quantities of Fe.




Figure 9: Performance of different acquisition functions in a 40-experiment budget AL run with 100 picoseconds MD simulation time. All functions start from identical initial sets. Open symbols represent MD simulated melting temperatures. The filled symbols with error bars represent RF predicted melting temperatures. Xx represents other elements, Cu and Cr. The insert includes a close-up on the best performing MPCAs compositions. These compositions contain high quantities of Fe.




Figure 10: Performance of different acquisition functions in a 40-experiment budget AL run with 200 picoseconds MD simulation time. All functions start from identical initial sets. Open symbols represent MD simulated melting temperatures. The filled symbols with error bars represent RF predicted melting temperatures. Xx represents other elements, Cu and Cr. The insert includes a close-up on the best performing MPCAs compositions. These compositions contain high quantities of Fe.

Table 2: *Best results from information acquisition functions running simulations for 50 picoseconds. RF prediction melting temperature and uncertainty taken when best composition selected for each acquisition function.*

ACQUISITION FUNCTION	MD SIMULATED T _M (K)	RF PREDICTED T _M (K)	RF UNCERTAINTY (K)	EXPERIMENT	COMPOSITION
MM	2472	2345	71	31	Fe 50%, Ni 50%
UCB	2498	2416	63	10	Co 10%, Fe 50%, Ni 40%
MLI	2502	2348	98	12	Fe 50%, Ni 50%
MEI	2492	2403	82	12	Co 20%, Fe 50%, Ni 30%
MU	2282	2137	118	-	Co 40%, Fe 20%, Ni 40%
RAND	2382	2202	52	-	Cu 10%, Fe 50%, Ni 40%

Table 3: *Best results from information acquisition functions running simulations for 100 picoseconds. RF prediction melting temperature and uncertainty taken when composition selected for each acquisition function.*

ACQUISITION FUNCTION	MD SIMULATED T _M (K)	RF PREDICTED T _M (K)	RF UNCERTAINTY (K)	EXPERIMENT	COMPOSITION
MM	2489	2460	36	24	Fe 50%, Ni 50%
UCB	2521	2422	78	15	Fe 50%, Ni 50%
MLI	2539	2455	60	15	Fe 50%, Ni 50%
MEI	2486	2415	87	12	Co 30%, Fe 50%, Ni 20%
MU	2299	2160	117	-	Co 20%, Cu 20%, Fe 50%, Ni 10%
RAND	2382	2202	52	-	Cu 10%, Fe 50%, Ni 40%

Table 4: *Best results from information acquisition functions running simulations for 200 picoseconds. RF prediction melting temperature and uncertainty taken when composition selected for each acquisition function.*

ACQUISITION FUNCTION	MD SIMULATED T _M (K)	RF PREDICTED T _M (K)	RF UNCERTAINTY (K)	EXPERIMENT	COMPOSITION
MM	2485	2461	29	27	Fe 50%, Ni 50%
UCB	2487	2417	81	17	Co 30%, Fe 50%, Ni 20%
MLI	2480	2413	85	14	Co 10%, Fe 50%, Ni 40%
MEI	2476	2418	91	9	Co 10%, Fe 50%, Ni 50%
MU	2285	2160	118	-	Co 40%, Cu 10%, Fe 30%, Ni 20%
RAND	2432	2360	80	-	Co 40%, Fe 40%, Ni 20%

Quite interestingly, all acquisition functions that take exploitation into consideration (MLI, UCB, MEI, and MM) find the alloys with the highest melting temperatures in few experiments, regardless of the simulation time used. This is quite remarkable given the significant variability in output for the 50 ps long simulations. We find that highest melting temperature alloys consist of 50 at.% Fe with the remaining 50 at.% distributed between Ni and Co; this set of alloys are predicted

to have nearly identical melting temperatures, see Figure 11. As expected, the MU function and random exploration do not explore high-melting temperature compositions within the 40 experiment budget. The graph seen for MM, UCB, MLI, and MEI has been made for MU and random and are in the SI as S3, S4. We find that RF prediction for the highest melting temperature alloy tends to underestimate the MD result, see Tables 2, 3 and 4, this is mostly true during the entire active learning workflows themselves, see Figures 8, 9 and 10. This can be attributed to the nature of the RF models, which are based on decision trees that saturate when extrapolating to values outside of their training.




Figure 11: Predicted melting temperature best performing alloys with composition $Fe_{50}Co_xNi_{50-x}$. Mean values and uncertainty estimates over 36 independent runs are shown.

In all our tests, we find that the acquisition functions that combine both exploration and exploitation (UCB, MLI, MEI) consistently outperform the purely exploitative function (MM), which can spend resources in local maxima. As seen in Tables 2-4 UCB, MLI, and MEI find the optimal composition within less than 17 experiments while MM found it past 27. While in our example MM finds the family of best performer candidates, it takes approximately twice as many experiments as the other functions.

3.2 Noisy data and RF uncertainty

We find that the AL workflow can find best performer alloys in approximately the same number of iterations regardless of the simulation time use. Between 9 and 17 simulations were required for the UCB, MEI, and MLI strategies. As mentioned above, this is despite the significant level of noise, especially in the 50 ps runs. We also find that while longer MD simulations do not result in a reduction in the number of experiments required to achieve the design goal, it results in improved accuracy of the RF prediction. Tables 2, 3 and 4 show that the workflow with 50 ps runs significantly underestimates the temperature of the optimal alloy and the uncertainty estimates are overly optimistic.

To better understand how the models gain knowledge in the presence of noise we explore the evolution of the model predictions for $Fe_{50}Co_xNi_{50-x}$ alloys, the top performers, at various stages

during the AL process. Figure 12 shows mean and uncertainty estimate predictions at four stages of an MLI workflow for the 50 ps case; results for 100 and 200 ps simulations show similar trends and are included in the SI as Figures S5, S6. Open circles denote MD simulations not already explored at the corresponding cycle and filled circles represent the same values for compositions that have been explored. We note that noise in the data affects both mean predictions and uncertainty estimates and can thus have non-trivial effects on AL decisions. Figure 12 shows that even with only the initial 39 datapoints (Iteration 0), the model predicts relatively high melting temperature for these alloys, comparable to the highest melting temperature materials in the initial set, see Figure 8. As the AL workflow with MLI explore compositions, the mean prediction for the selected family improves from 12 (a) to (b) as the mean moves towards the true MD temperatures of the alloys. At step 12, Fig. 12 (d), the model predicts high melting temperatures for the alloys with $x=0$, 10, 20 % and $x=0$ is selected for simulation. In subsequent steps, the MLI algorithm hones in the rest of the family and improves the model accuracy. When comparing the results for 50 ps runs with those of 100 and 200 ps (SI Figures S5 and S6) we observe similar paths to the optimal alloys with a decrease in uncertainties for longer simulation times. This is because the variability in the MD simulations is relatively small compared to the differences in melting temperatures across alloys.




Figure 12: Red lines represent the predicted mean melting temperature and shaded region represents the uncertainty estimates for predictions on compositions as a function of Co content in $Fe_{50}Co_xNi_{50-x}$ at various stages of the AL workflow. Whereby x starts at 0% and goes to 50% with 1% step size. Results correspond to the MLI function with a 50 ps MD simulation time. Open circles represent MD-simulated values unknown to the model at the time and filled symbols represent the values included in the model.

4 Conclusions

The combination of machine learning and physics-based simulations holds great promise to guide the experimental search and accelerate the development of novel materials. The use of surrogate-based models that can make use of computational simulations instead of experiments can save resources and time in design cycles. We find the molecular dynamics simulations, even under significant stochastic noise, can be a powerful tool to be used in active learning efforts.

We developed an AL workflow with MD simulations for the discovery of MPCAs materials for high-temperature applications. Starting from a limited initial set, we created a set of descriptors and a RF model to use as a surrogate for the optimization of our AL model. We determined that RF models are capable of handling multi-dimensional spaces with a limited initial dataset. We paired our RF model with a cloud-based simulation environment for automatic querying of melting temperatures from MD simulations with an interatomic potential for FCC alloys. We analyzed the effect of MD simulation time in the overall uncertainty quantification when running such calculations

and showed that it has an important effect in capturing the sample-to-sample variability caused by the stochastic nature of MD. We compared the performance of four different acquisition functions in a closed, but high-dimensional, search space. We found that the best performing acquisition functions, MLI, MEI, and UCB, are combinations of exploitation and exploration.

The workflow we developed can be adapted toward the optimization of other material properties and to the use of other sources of information other than computational simulations. We provide the analysis and all other relevant information in a public SimTool on nanoHUB.org.

5 Models and Data Availability

The entire simulation workflow to calculate melting temperature, including pre- and post-analysis was implemented as a SimTool⁸ in nanoHUB⁷⁰ and is available for online simulation.⁷¹ SimTools⁸ are complete simulation workflows with verified inputs and outputs that enable workflows and the data they generate to be findable, accessible, and reproducible. An additional advantage of using SimTools is that every simulation performed in nanoHUB is cached, so that given identical input parameters, cached results are retrieved rather than re-running a simulation. The stand-alone SimTool⁷¹ and the AL workflow are available for online simulations in nanoHUB.⁷²

Acknowledgements

This effort was supported by the US National Science Foundation (DMREF-1922316). We acknowledge computational resources from nanoHUB and Purdue University through the Network for Computational Nanotechnology. J. C. V. thanks the Science and Technology Council of México (Consejo Nacional de Ciencia y Tecnología, CONACYT) for partial financial support of this research.

Ethics declarations

Conflict of interest

On behalf of all authors, the corresponding author states that there is no conflict of interest.

Author Information

Corresponding Author

Alejandro Strachan - School of Materials Engineering and Birck Nanotechnology Center, Purdue University, West Lafayette, Indiana 47907; Email: strachan@purdue.edu

References

- [1] Natalie S Eyke, William H Green, and Klavs F Jensen. Iterative experimental design based on active machine learning reduces the experimental burden associated with reaction screening. *Reaction Chemistry & Engineering*, 5(10):1963–1972, 2020.

- [2] Lifei Zhao, Zhen Li, Bruce Caswell, Jie Ouyang, and George Em Karniadakis. Active learning of constitutive relation from mesoscopic dynamics for macroscopic modeling of non-newtonian flows. *Journal of Computational Physics*, 363:116–127, 2018.
- [3] A Gilad Kusne, Heshan Yu, Changming Wu, Huairuo Zhang, Jason Hattrick-Simpers, Brian DeCost, Suchismita Sarker, Corey Oses, Cormac Toher, Stefano Curtarolo, et al. On-the-fly closed-loop materials discovery via bayesian active learning. *Nature communications*, 11(1):1–11, 2020.
- [4] Dezhen Xue, Prasanna V Balachandran, John Hogden, James Theiler, Deqing Xue, and Turab Lookman. Accelerated search for materials with targeted properties by adaptive design. *Nature communications*, 7(1):1–9, 2016.
- [5] Kevin Tran and Zachary W Ulissi. Active learning across intermetallics to guide discovery of electrocatalysts for co2 reduction and h2 evolution. *Nature Catalysis*, 1(9):696–703, 2018.
- [6] Anubhav Jain, Shyue Ping Ong, Geoffroy Hautier, Wei Chen, William Davidson Richards, Stephen Dacek, Shreyas Cholia, Dan Gunter, David Skinner, Gerbrand Ceder, et al. Commentary: The materials project: A materials genome approach to accelerating materials innovation. *APL materials*, 1(1):011002, 2013.
- [7] Atsuto Seko, Atsushi Togo, Hiroyuki Hayashi, Koji Tsuda, Laurent Chaput, and Isao Tanaka. Prediction of low-thermal-conductivity compounds with first-principles anharmonic lattice-dynamics calculations and bayesian optimization. *Physical review letters*, 115(20):205901, 2015.
- [8] Saaketh Desai, Steven Clark, and Alejandro Strachan. Introduction to simtools, Apr 2020. URL: <https://nanohub.org/resources/introtosimtools>, DOI: 10.21981/4Z7T-X415.
- [9] Prasanna V Balachandran, Dezhen Xue, James Theiler, John Hogden, and Turab Lookman. Adaptive strategies for materials design using uncertainties. *Scientific reports*, 6(1):1–9, 2016.
- [10] Chiho Kim, Anand Chandrasekaran, Anurag Jha, and Rampi Ramprasad. Active-learning and materials design: The example of high glass transition temperature polymers. *MRS Communications*, 9(3):860–866, 2019.
- [11] Julia Ling, Maxwell Hutchinson, Erin Antono, Sean Paradiso, and Bryce Meredig. High-dimensional materials and process optimization using data-driven experimental design with well-calibrated uncertainty estimates. *Integrating Materials and Manufacturing Innovation*, 6(3):207–217, 2017.
- [12] Juan C Verduzco, Ernesto E Marinero, and Alejandro Strachan. An active learning approach for the design of doped LLZO ceramic garnets for battery applications. *Integrating Materials and Manufacturing Innovation*, pages 1–12, 2021.
- [13] National Science and Technology Council (US). *Materials genome initiative for global competitiveness*. Executive Office of the President, National Science and Technology Council, 2011.
- [14] John Wise, Alexandra Grebe de Barron, Andrea Splendiani, Beeta Balali-Mood, Drashtti Vasant, Eric Little, Gaspare Mellino, Ian Harrow, Ian Smith, Jan Taubert, et al. Implementation and relevance of fair data principles in biopharmaceutical r&d. *Drug discovery today*, 24(4):933–938, 2019.

- [15] Burr Settles. *Synthesis Lectures on Artificial Intelligence and Machine Learning: Active Learning*. Morgan & Claypool Publishers, San Rafael, United States, 2011.
- [16] Yanyao Shen, Hyokun Yun, Zachary C Lipton, Yakov Kronrod, and Animashree Anandkumar. Deep active learning for named entity recognition. *arXiv preprint arXiv:1707.05928*, 2017.
- [17] Pilsun Yoo, Michael Sakano, Saaketh Desai, Md Mahbubul Islam, Peilin Liao, and Alejandro Strachan. Neural network reactive force field for C, H, N, and O systems. *npj Computational Materials*, 7(1):1–10, 2021.
- [18] Aaron Gilad Kusne, Tieren Gao, Apurva Mehta, Liqin Ke, Manh Cuong Nguyen, Kai-Ming Ho, Vladimir Antropov, Cai-Zhuang Wang, Matthew J Kramer, Christian Long, et al. On-the-fly machine-learning for high-throughput experiments: search for rare-earth-free permanent magnets. *Scientific reports*, 4(1):1–7, 2014.
- [19] Pavel Nikolaev, Daylond Hooper, Frederick Webber, Rahul Rao, Kevin Decker, Michael Krein, Jason Poleski, Rick Barto, and Benji Maruyama. Autonomy in materials research: a case study in carbon nanotube growth. *npj Computational Materials*, 2(1):1–6, 2016.
- [20] Fang Ren, Logan Ward, Travis Williams, Kevin J Laws, Christopher Wolverton, Jason Hattrick-Simpers, and Apurva Mehta. Accelerated discovery of metallic glasses through iteration of machine learning and high-throughput experiments. *Science advances*, 4(4):eaq1566, 2018.
- [21] Marcus M Noack, Kevin G Yager, Masafumi Fukuto, Gregory S Doerk, Ruipeng Li, and James A Sethian. A kriging-based approach to autonomous experimentation with applications to x-ray scattering. *Scientific reports*, 9(1):1–19, 2019.
- [22] Shin Kiyohara, Hiromi Oda, Koji Tsuda, and Teruyasu Mizoguchi. Acceleration of stable interface structure searching using a kriging approach. *Japanese Journal of Applied Physics*, 55(4):045502, 2016.
- [23] Erik Van Der Giessen, Peter A Schultz, Nicolas Bertin, Vasily V Bulatov, Wei Cai, Gábor Csányi, Stephen M Foiles, Marc GD Geers, Carlos González, Markus Hütter, et al. Roadmap on multiscale materials modeling. *Modelling and Simulation in Materials Science and Engineering*, 28(4):043001, 2020.
- [24] Alberto M Cuitiño, Laurent Stainier, Guofeng Wang, Alejandro Strachan, Tahir Çağın, William A Goddard, and Michael Ortiz. A multiscale approach for modeling crystalline solids. *Journal of computer-aided materials design*, 8(2):127–149, 2001.
- [25] Alejandro Strachan, Tahir Çağın, and William A Goddard III. Phase diagram of mgo from density-functional theory and molecular-dynamics simulations. *Physical Review B*, 60(22):15084, 1999.
- [26] Adri CT Van Duin, Siddharth Dasgupta, Francois Lorant, and William A Goddard. Reaxff: a reactive force field for hydrocarbons. *The Journal of Physical Chemistry A*, 105(41):9396–9409, 2001.
- [27] Albert P Bartók, Mike C Payne, Risi Kondor, and Gábor Csányi. Gaussian approximation potentials: The accuracy of quantum mechanics, without the electrons. *Physical review letters*, 104(13):136403, 2010.

- [28] Aidan P Thompson, Laura P Swiler, Christian R Trott, Stephen M Foiles, and Garritt J Tucker. Spectral neighbor analysis method for automated generation of quantum-accurate interatomic potentials. *Journal of Computational Physics*, 285:316–330, 2015.
- [29] Jörg Behler and Michele Parrinello. Generalized neural-network representation of high-dimensional potential-energy surfaces. *Physical review letters*, 98(14):146401, 2007.
- [30] Thaer M. Dieb, Shenghong Ju, Kazuki Yoshizoe, Zhufeng Hou, Junichiro Shiomi, and Koji Tsuda. Mmts: automatic complex materials design using monte carlo tree search. *Science and technology of advanced materials*, 18(1):498–503, 2017.
- [31] Thaer M Dieb, Shenghong Ju, Junichiro Shiomi, and Koji Tsuda. Monte carlo tree search for materials design and discovery. *MRS Communications*, 9(2):532–536, 2019.
- [32] Tarak K Patra, Troy D Loeffler, and Subramanian KRS Sankaranarayanan. Accelerating copolymer inverse design using monte carlo tree search. *Nanoscale*, 12(46):23653–23662, 2020.
- [33] Troy D Loeffler, Suvo Banik, Tarak K Patra, Michael Sternberg, and Subramanian KRS Sankaranarayanan. Reinforcement learning in discrete action space applied to inverse defect design. *Journal of Physics Communications*, 5(3):031001, 2021.
- [34] Stéphane Gorsse, Jean-Philippe Couzinié, and Daniel B Miracle. From high-entropy alloys to complex concentrated alloys. *Comptes Rendus Physique*, 19(8):721–736, 2018.
- [35] Yiping Lu, Yong Dong, Sheng Guo, Li Jiang, Huijun Kang, Tongmin Wang, Bin Wen, Zhijun Wang, Jinchuan Jie, Zhiqiang Cao, et al. A promising new class of high-temperature alloys: eutectic high-entropy alloys. *Scientific reports*, 4(1):1–5, 2014.
- [36] ZQ Xu, ZL Ma, M Wang, YW Chen, YD Tan, and XW Cheng. Design of novel low-density refractory high entropy alloys for high-temperature applications. *Materials Science and Engineering: A*, 755:318–322, 2019.
- [37] Sathiyamoorthi Praveen and Hyoung Seop Kim. High-entropy alloys: Potential candidates for high-temperature applications—an overview. *Advanced Engineering Materials*, 20(1):1700645, 2018.
- [38] O.N. Senkov, G.B. Wilks, D.B. Miracle, C.P. Chuang, and P.K. Liaw. Refractory high-entropy alloys. *Intermetallics*, 18(9):1758–1765, 2010.
- [39] WL Wang, L Hu, SB Luo, LJ Meng, DL Geng, and B Wei. Liquid phase separation and rapid dendritic growth of high-entropy CoCrCuFeNi alloy. *Intermetallics*, 77:41–45, 2016.
- [40] Feng He, Zhijun Wang, Qingfeng Wu, Junjie Li, Jincheng Wang, and CT Liu. Phase separation of metastable CoCrFeNi high entropy alloy at intermediate temperatures. *Scripta Materialia*, 126:15–19, 2017.
- [41] Qi-Jun Hong, Jan Schroers, Douglas Hofmann, Stefano Curtarolo, Mark Asta, and Axel van de Walle. Theoretical prediction of high melting temperature for a Mo–Ru–Ta–W HCP multi-principal element alloy. *npj Computational Materials*, 7(1):1–4, 2021.
- [42] James R Morris, CZ Wang, KM Ho, and CT Chan. Melting line of aluminum from simulations of coexisting phases. *Physical Review B*, 49(5):3109, 1994.

- [43] James R Morris and Xueyu Song. The melting lines of model systems calculated from coexistence simulations. *The Journal of chemical physics*, 116(21):9352–9358, 2002.
- [44] Diana Farkas and Alfredo Caro. Model interatomic potentials and lattice strain in a high-entropy alloy. *Journal of Materials Research*, 33(19):3218–3225, 2018.
- [45] Sheng-Nian Luo, Thomas J Ahrens, Tahir Çağın, Alejandro Strachan, William A Goddard III, and Damian C Swift. Maximum superheating and undercooling: Systematics, molecular dynamics simulations, and dynamic experiments. *Physical Review B*, 68(13):134206, 2003.
- [46] Diana Farkas and Alfredo Caro. Model interatomic potentials for fe–ni–cr–co–al high-entropy alloys. *Journal of Materials Research*, 35(22):3031–3040, 2020.
- [47] A. S. Lemak and N. K. Balabaev. On The Berendsen Thermostat. *Molecular Simulation*, 13(3):177–187, 1994.
- [48] Peter J. Daivis and Denis J. Evans. Comparison of constant pressure and constant volume nonequilibrium simulations of sheared model decane. *J. Chem. Phys*, 100, 1994.
- [49] Peter Mahler Larsen, Søren Schmidt, and Jakob Schiøtz. Robust structural identification via polyhedral template matching. *Modelling and Simulation in Materials Science and Engineering*, 24(5):055007, 2016.
- [50] Alexander Stukowski. Visualization and analysis of atomistic simulation data with OVITO—the open visualization tool. *Modelling and Simulation in Materials Science and Engineering*, 18(1):015012, 2009.
- [51] Xiang-Guo Li, Chi Chen, Hui Zheng, Yunxing Zuo, and Shyue Ping Ong. Complex strengthening mechanisms in the nbmotaw multi-principal element alloy. *npj Computational Materials*, 6(1):1–10, 2020.
- [52] Sheng Yin, Yunxing Zuo, Anas Abu-Odeh, Hui Zheng, Xiang-Guo Li, Jun Ding, Shyue Ping Ong, Mark Asta, and Robert O Ritchie. Atomistic simulations of dislocation mobility in refractory high-entropy alloys and the effect of chemical short-range order. *Nature communications*, 12(1):1–14, 2021.
- [53] Tatiana Kostiuchenko, Fritz Körmann, Jörg Neugebauer, and Alexander Shapeev. Impact of lattice relaxations on phase transitions in a high-entropy alloy studied by machine-learning potentials. *npj Computational Materials*, 5(1):1–7, 2019.
- [54] Qi-Jun Hong and Axel van de Walle. Solid-liquid coexistence in small systems: A statistical method to calculate melting temperatures. *The Journal of chemical physics*, 139(9):094114, 2013.
- [55] Z. Wu, H. Bei, F. Otto, G.M. Pharr, and E.P. George. Recovery, recrystallization, grain growth and phase stability of a family of fcc-structured multi-component equiatomic solid solution alloys. *Intermetallics*, 46:131–140, 2014.
- [56] Juan J Valencia and Peter N Quested. Thermophysical properties. *NIST*, 2013. Table 5: Thermal expansion of selected pure metals at temperatures close to melting.

- [57] Rohit K Tripathy and Ilias Bilonis. Deep UQ: Learning deep neural network surrogate models for high dimensional uncertainty quantification. *Journal of computational physics*, 375:565–588, 2018.
- [58] Tianqi Chen and Carlos Guestrin. Xgboost: A scalable tree boosting system. In *Proceedings of the 22nd ACM sigkdd international conference on knowledge discovery and data mining*, pages 785–794, 2016.
- [59] Leo Breiman. Random forests. *Machine learning*, 45(1):5–32, 2001.
- [60] Bradley Efron. *Model Selection Estimation and Bootstrap Smoothing*. Division of Biostatistics, Stanford University, 2012.
- [61] Stefan Wager, Trevor Hastie, and Bradley Efron. Confidence intervals for random forests: The jackknife and the infinitesimal jackknife. *The Journal of Machine Learning Research*, 15(1):1625–1651, 2014.
- [62] M Hutchinson. Python package: Lolo, 2016. URL: <https://github.com/CitrineInformatics/lolo>.
- [63] M Hutchinson. Python package: lolopy, 2020. Version 1.1.1. URL: <https://pypi.org/project/lolopy/1.1.1/>.
- [64] Dipendra Jha, Logan Ward, Arindam Paul, Wei-keng Liao, Alok Choudhary, Chris Wolverton, and Ankit Agrawal. Elemnet: Deep learning the chemistry of materials from only elemental composition. *Scientific reports*, 8(1):1–13, 2018.
- [65] Yan Zhang, Cheng Wen, Changxin Wang, Stoichko Antonov, Dezhen Xue, Yang Bai, and Yanjing Su. Phase prediction in high entropy alloys with a rational selection of materials descriptors and machine learning models. *Acta Materialia*, 185:528–539, 2020.
- [66] Ziqing Zhou, Yeju Zhou, Quanfeng He, Zhaoyi Ding, Fucheng Li, and Yong Yang. Machine learning guided appraisal and exploration of phase design for high entropy alloys. *npj Computational Materials*, 5(1):1–9, 2019.
- [67] Shyue Ping Ong, William Davidson Richards, Anubhav Jain, Geoffroy Hautier, Michael Kocher, Shreyas Cholia, Dan Gunter, Vincent L Chevrier, Kristin A Persson, and Gerbrand Ceder. Python materials genomics (pymatgen): A robust, open-source python library for materials analysis. *Computational Materials Science*, 68:314–319, 2013.
- [68] Donald R Jones, Matthias Schonlau, and William J Welch. Efficient global optimization of expensive black-box functions. *Journal of Global optimization*, 13(4):455–492, 1998.
- [69] Emmanuel Vazquez and Julien Bect. Convergence properties of the expected improvement algorithm with fixed mean and covariance functions. *Journal of Statistical Planning and inference*, 140(11):3088–3095, 2010.
- [70] Alejandro Strachan, Gerhard Klimeck, and Mark Lundstrom. Cyber-enabled simulations in nanoscale science and engineering. *Computing in Science & Engineering*, 12(2):12–17, 2010.

- [71] Zachary D McClure, Saaketh Desai, and Alejandro Strachan. High entropy alloy melting point calculation, Mar 2020. URL: <https://nanohub.org/resources/meltheas>, DOI: 10.21981/W5VD-T039.
- [72] Juan Carlos Verduzco Gastelum, David Enrique Farache, Zachary D McClure, Saaketh Desai, and Alejandro Strachan. Active learning workflow for mpcas, Oct 2021. URL: <https://nanohub.org/resources/activemeltheas>, DOI: 10.21981/NK7E-HA16.