

Bayesian blacksmithing: Discovering thermomechanical properties and deformation mechanisms in high-entropy refractory alloys

Jacob Startt¹, Megan J. McCarthy², Mitchell A. Wood², Sean Donegan³, and Rémi Dingreville^{1,*}

¹Center for Integrated Nanotechnologies, Sandia National Laboratories, Albuquerque, NM 87185, USA

²Center for Computing Research, Computational Multiscale Department, Sandia National Laboratories, Albuquerque, NM, 87185, USA

³Manufacturing and Industrial Technologies Division, Air Force Research Laboratory , Wright Patterson AFB, OH 45433, USA

*rdingre@sandia.gov

ABSTRACT

Finding alloys with specific design properties is challenging due to the large number of possible compositions and the complex interactions between elements. This study introduces a multi-objective Bayesian optimization approach guiding molecular dynamics simulations for discovering high-performance refractory alloys with both targeted intrinsic static thermomechanical properties and also deformation mechanisms occurring during dynamic loading. The objective functions are aiming for excellent thermomechanical stability via a high bulk modulus, a low thermal expansion, a high heat capacity, and for a resilient deformation mechanism maximizing the retention of the BCC phase after shock. Contrasting two optimization procedures, we show that the Pareto-optimal solutions are confined to a small performance space when the property objectives display a cooperative relationship. Conversely, the Pareto front is much broader in the performance space when these properties have antagonistic relationships. Density functional theory simulations validate these findings and unveil underlying atomic-bond changes driving property improvements.

INTRODUCTION

Materials discovery and optimization are at the heart of technological advancement and innovation. The largest and most ambitious engineering projects today depend on the discovery and delivery of new materials. Examples include fusion energy which require materials capable of withstanding the harsh conditions of plasma confinement,¹ next-generation aerospace applications needing advanced materials for propulsion systems, high radiation environments, and re-entry events,² or quantum computing, which would greatly benefit from new quantum and superconducting materials.³

Human-guided formulations of novel materials solutions can be an impediment to these innovations, hindered by time-consuming and costly experimental steps and inefficient exploration procedures governed by Edisonian approaches.^{4,5} In particular, over the past decade, the rapid emergence of multi-principal element alloys (MPEA), high-entropy alloys (HEA) and complex, concentrated alloys (CCA)⁶ has highlighted the almost infinite number (10^{117} according to Cantor⁷) of alloys awaiting to be explored. Because of the size and complexity of the alloy composition space, research on MPEA and HEA materials has typically been limited to alloys at equiatomic or near-equiatomic compositions.^{6,8–11} Advances in high-throughput characterization and computing techniques have gradually enabled investigations to peer beyond the equiatomic ratios and explore the vast CCA compositional space.^{8,12–16} For example, in a previous study, we utilized high-throughput density functional theory (DFT) calculations to show that higher concentrations of Mo in the refractory MoNbTaTi HEA led to improvement in mechanical properties including elastic stiffness and yield strength as a result of a covalent-like strengthening from the Mo bonds.¹⁷ Similarly, Zhang and coworkers¹⁸ used high-throughput DFT simulations to screen over 20 000 compositions to predict the phase evolution in NbVZr-Ti_x refractory CCAs (RCCAs). In another example, Shin and coworkers¹⁹ used a high-throughput CALPHAD thermodynamic approach combined with experimental creep datasets to predict the creep behavior of high-temperature alumina-forming austenitic stainless-steel alloys.

While the aforementioned examples clearly demonstrate the usefulness of compositional tuning in MPEA, HEAs, and CCAs, they each share a common theme in that the “exploration” of the design space (*i.e.*, the composition) was performed on a limited compositional space and a human derived grid-like search pattern. In an example of this approach, one element’s




Figure 1. Multi-objective Bayesian optimization workflow. The multi-objective function consists of optimizing both thermomechanical properties and target a specific deformation mechanism in MoNbTaTi alloys. The initial database of properties is based on molecular dynamics (MD) simulations on a handful of compositions. The multi-objective Bayesian optimization is based on the Thompson sampling efficient multi-objective optimization (TS-EMO) algorithm.²⁷

concentration is typically varied at a time while the other element concentrations are changed according to some simplistic rule or algorithm, and while this approach can yield desirable results, the exploration of the full compositional space is limited and offers no guarantee of finding an optimal property state. Conversely, Bayesian optimization techniques are explicitly designed to achieve optimization within the entire performance space, and since they generally require a small number of evaluations to achieve the desired optimization, they are ideally suited to the needs of alloy optimization problems. In Bayesian optimization, known or existing information is used to construct a prior distribution or probability model (often a Gaussian process), connecting the design features to objective properties, to find the set of Pareto-optimal solutions (*i.e.*, where no improvement in any objective can be achieved without simultaneously worsening at least one other objective). In alloy design problems, the Pareto-optimal solutions represent a subset of alloy compositions that optimize targeted material properties.²⁰ And so recently, Bayesian optimization and other machine-learning based optimization techniques have been applied to the discovery of novel alloys to improve various material properties.^{21,22} Typically, in most of these studies, alloy properties were estimated using a quick-to-evaluate acquisition function based on an analytical model, like Vegard's law or the Maresca-Curtin model for yield strength in BCC HEAs.²³ Alternatively, DFT was also used albeit with a much higher computational cost. For instance, Pedersen and coworkers²⁴ used DFT-based Bayesian optimization to predict the most active compositions for the electrochemical oxygen reduction reaction in AgIrPdPtRu and IrPdPtRuRh HEAs. In another example, Khatamsaz and coworkers^{25,26} employed a multi-objective Bayesian optimization scheme using the Maresca-Curtin analytical model in conjunction with a constraint screened design space via CALPHAD and Thermo-Calc estimations to simultaneously maximize the yield strength, the Pugh ratio (ductility/brittleness), and the Cauchy pressure in high-temperature refractory alloys.

In this paper, we extend this body of work to go beyond the prior optimizations of static alloy properties by targeting an optimized response to dynamic loading, in which the mechanisms of deformation themselves are optimized alongside other (static) thermomechanical properties. Determination of dynamic properties and behaviors requires more than an analytical model such as the Maresca-Curtin model necessitating us to use instead atomistic simulations as an acquisition function in our optimization workflow to explicitly simulate those deformation mechanisms. Specifically, we utilized molecular dynamics (MD) simulations to calculate each targeted property and achieve more accurate predictions of the mechanical performance and deformation mechanisms under dynamic loading which can often be underestimated²⁸ or unable to be predicted with an analytical model such as the Maresca-Curtin model. Molecular dynamics offers two unique advantage of (i) enabling optimization schemes that are outside the time and length scale scopes of both DFT simulations²⁴ and experiments²⁶ and (ii) the resulting optimizations are driven directly by the dynamic, physical mechanisms underlying property changes. With this atomistic-based optimization approach, our goal is to identify optimal compositions in the refractory MoNbTaTi CCA that yield high bulk moduli (B), low coefficients of thermal expansion (α_L), and high heat capacities (C_P) as static properties, and for which the fraction of BCC phase retained after shock loading (hereby referred to as BCC%) is maximized as the preferred deformation mechanism as opposed to the crystalline-to-amorphous phase transformation classically seen in many shocked alloys.²⁹⁻³⁴ The static thermomechanical properties are selected to discover alloys with excellent and improved




Figure 2. Performance space optimization. Three-dimensional projections of the MD calculated property performance spaces of optimizations A (panel a) and B (panel b). Pareto optimal points are colored purple while points from the initial database are colored turquoise. Points evaluated during optimization but which do not lie on either Pareto front are colored grey. To aid in visualization, two-dimensional projections are drawn onto the visible faces, and shadows are projected onto the points within the three-dimensional volume.

high-temperature thermodynamic properties, while the mechanism associated with the restoration of the BCC phase upon shock is chosen as a resilient deformation mechanism. We performed two computationally driven multi-objective Bayesian optimizations each with four design variables (dimension of the compositional space) and three alloy performance objectives (two thermomechanical properties and the BCC retention deformation mechanism). For both optimizations, the multi-objective optimization consisted of maximizing both B and BCC% as two of the three performance objectives. The two optimizations differ in the selection of the third objective, where in the first optimization, we seek to minimize α_L , while in the second we seek to maximize C_p . We employed the Thompson sampling efficient multi-objective optimization (TS-EMO)²⁷ algorithm in which surrogate objective functions trained on molecular dynamics data are used to iteratively sample the input compositional space and predict new compositions with improved alloy performance objectives. A high-level scheme of the TS-EMO algorithm and process flow is presented Fig. 1. Details of the Bayesian optimization framework (including the TS-EMO algorithm) and atomistic calculations are provided in the Methods section. Contrasting both optimization results, we find that B , BCC%, and α_L are synergistic to one another and optimize to a narrow compositional space with compositions that are heavy in Mo and Ta, and light in Ti. Conversely, C_p optimizes antagonistically to B and BCC%, creating an extended Pareto front from regions of high strength which are rich in Mo and Ta but low in Ti to regions of moderate strength coupled with high C_p which have decreased Ta concentrations in favor of increased Ti and Nb concentrations. Confirmatory DFT simulations and an analysis of property correlation reveal the origin of improved properties in terms of changes in the electronic structure, altering the nature of atomic bonds. By optimizing both static thermomechanical properties and dynamic deformation mechanisms, this approach not only serves as a tool for discovering optimal compositions with targeted properties (as it is classically done in generic Bayesian optimization frameworks), but also as a mean to discover novel mechanisms and reveal structure-property relationships in the compositional space. Indeed, the addition of a mechanism as one of the objective functions enables the navigation of the Pareto front to hone in on specific compositions to discover peculiar and targeted mechanisms which can be investigated later on in greater details to gain understanding on their origin.

RESULTS

Pareto-optimal Solutions

We have performed a pair of alloy optimizations, hereby denoted ‘optimization A’ and ‘optimization B’, where the goal in each optimization was to simultaneously optimize three different material properties. In optimization A, α_L was minimized, while

B and $BCC\%$ were maximized as a function of alloy composition, whereas in optimization B, all three objectives, C_p , B and $BCC\%$ were maximized again as a function of composition. Examples of molecular dynamics simulations used to calculate each objective property along with more detailed descriptions of the objective calculations are presented in the Methods section. To ensure convergence of the Pareto front, each optimization was carried out for 100 cycle evaluations, evaluating one new composition per cycle (*i.e.*, a batch size equal to 1), repeating each full optimization starting from the same initial dataset for a total of three times each.

The optimized performance spaces of optimizations A and B are shown in Figure 2, where the relationships between each of the three material properties are mapped to one-another in a three-dimensional space. The individual points in these plots represent single alloy compositions, evaluated either as part of the initial dataset (turquoise) or during the optimizations (purple and grey), while the positions of the points in the XYZ space indicate the MD computed values of each material property (X = coefficient of thermal expansion for optimization A or X = heat capacity for optimization B; Y = bulk modulus for both optimizations; Z = retained BCC for both optimizations). Alloy compositions that were found to lie on the Pareto front are colored in purple, while non-Pareto compositions (*i.e.*, evaluated during the optimizations but that did not lie on the Pareto front) are colored in grey. Upon inspection, it is immediately observable that the two optimizations exhibit starkly contrasting behaviors, with each converging to very different Pareto optimal regions. In optimization A (Fig. 2a), the Pareto front was contained to a small localized region in the corner of the performance space. In this case, the three objectives of (α_L , B , and $BCC\%$) were all collectively optimized by the same compositions, indicating a cooperative relationship shared between the three properties with respect to composition. This relationship is understandable, as each of these properties represents a measure of the mechanical deformation of atomic bonds and lattice structure in response to some thermodynamic force. In the case of the bulk modulus, B , that force is simply pressure and the modulus is measured as a function of the volumetric stress or compression of the lattice. Likewise, α_L is a measure of the volumetric changes in the lattice in response to changes in temperature. Lastly, $BCC\%$ is a cumulative measure of the reversible deformations that potentially occur in response to both pressure and temperature changes that are both dynamic and extreme. Thus, it is a reasonable expectation that optimization of one of these properties would proportionally affect the others.

In optimization B (Fig. 2b), the Pareto region exhibits a more complex shape compared with the localized Pareto region of optimization A. The Pareto front is now drawn out across the performance space, forming a narrow arc that stretches from the upper corner, where alloys have excellent mechanical performance, high stiffness, and low heat capacity, to the opposite lower corner where alloys exhibit poor mechanical strength and stiffness but excellent heat capacity. In this case, the three objectives (C_p , B , and $BCC\%$) display antagonistic relationships as seen in the two-dimensional projections of the performance space. While B and $BCC\%$ go in tandem through their dependence on volumetric deformation, the third objective, C_p , sees an opposite trend as it is no longer associated directly with a volumetric response in the lattice, instead being directly proportional to changes in internal energy and temperature (see Eq. (2) in Methods). Thus, it is reasonable to conclude that the heat capacity exhibits a compositional decoupling with the other material properties of interest. In fact, the relationship observed between the heat capacity and the bulk modulus goes beyond that of just a simple decoupling. In our previous DFT-based investigation of the MoNbTaTi system,¹⁷ we showed that the bulk modulus and heat capacity shared an inversely proportional relationship to one another. We observe the same correlation between the two properties here in optimization B as evidenced by the nearly linear shape of the Pareto front in the $B - C_p$ projection of Fig. 2b.

While the coefficient of expansion (α_L) and heat capacity (C_p) were never optimized together in this work, we can provide some estimation of their relationship in the MoNbTaTi alloy. In Fig. 3, we show two-dimensional projections of the performance space for optimization B, depicting the bulk modulus (B) and the fraction of BCC phase retained ($BCC\%$) each plotted against C_p . In this case however, the color of each point represents the value of α_L , calculated separately through MD for each composition. Interestingly, the compositions that exhibit the lowest thermal expansions are most densely concentrated in the high $BCC\%$ retention and high stiffness regions of the performance space. Compositions that exhibit excellent heat capacity typically have poor thermal expansions, however, suggesting that these two properties are likely inversely correlated (however, we can not explicitly say this based on this data alone). Of note, α_L appears much more strongly correlated with B as opposed to $BCC\%$, as several compositions have low thermal expansions paired with poor BCC retention's (see dark points in the bottom right of Fig. 3b). The strong relationship with B further supports the notion these two properties are tied through thermodynamic volumetric relationships.

Composition Dependence of Pareto Optimal Solutions

We now turn our analysis towards the compositional dependence of the Pareto optimal solutions. Given the localization of the Pareto front of optimization A in a small performance space, we focus our attention to the more spread Pareto front of optimization B. In Fig. 4, we break down the Pareto front of optimization B into four regions, each of which characterizes a unique subset of optimal material properties. To infer compositional effects on the variation of properties, the distribution of the atomic species concentrations in each of the four regions are drawn as violin distribution plots. From here forward and for




Figure 3. Thermal expansion in optimization B Thermal expansion, α_L , was determined separately for every alloy composition evaluated in optimization B, and plotted as a function of color in **a** showing the bulk modulus, B , plotted against the heat capacity, C_P , and in **b**, showing the BCC retention plotted against C_P . The points with small black dots in the center represent the alloy composition in the Pareto front of optimization B. It is important to reiterate that the minimization of α_L was not included as an objective for optimization B, and its evaluation here is meant only to provide some understanding of the relationship between α_L with C_P .




Figure 4. Compositional dependence of Pareto optimal solutions. The Pareto region of optimization B is broken into four distinct regions, as denoted by the left side 3D plots, for which typical composition distributions are shown in the right side violin plots. Each Pareto region exhibits a unique set of property behaviors. **a** (B-I) Alloys are mechanically stiff and resilient to dynamic loading but poor capacity for heat absorption; **b** (B-II) mechanically resilient to dynamic loading but elastic stiffness is reduced in favor of improved heat absorption; **c** (B-III) slightly greater heat capacity at the cost of some stiffness and a significant reduction in mechanical resilience to dynamic loading; and **d** (B-IV) elastically ductile with excellent heat absorption capability but poor mechanical resiliency under dynamic loading. It should be noted that the Pareto points within region (B-I) are also representative of the localized Pareto optimal region found in optimization A, and exhibit nearly identical compositional dependencies. The arrows on the axes of the 3D performance space of region B-I indicate the direction of property optimization.

brevity, the individual regions will be referred to as regions B-I, B-II, B-III and B-IV as they are listed in Fig. 4. Region B-I corresponds to alloys exhibiting the highest B moduli and best BCC% retentions but with a low heat capacity. This region is also identical to the Pareto region of optimization A. Conversely, region IV corresponds to alloys with a high heat capacity but low bulk modulus. Regions B-II and B-III correspond to alloys with intermediate properties, transitioning from B-I to B-IV on the Pareto front. Looking at the compositional make up of alloys in region B-I in panel **a** of Fig. 4, we observe that these alloys are high in Mo and Ta and low in Ti. Moving on to region B-II, we note that the Ta concentrations in these alloys have become reduced in favor of an increased concentration of Nb. This change in compositional makeup has the effect of elastically softening the alloy (*i.e.* reducing B) while improving the heat capacity. The reduction in Ta concentration is most important for this shift, as Ta possesses the lowest heat capacity of the four metals in pure form³⁵ by a significant margin, and so by simple rule of mixture, a reduction of its concentration is necessary for improved heat capacitance. Interestingly, despite the decrease in B , the strength and resiliency of these alloys under dynamic loading remains fairly high (*i.e.* BCC% remains high), and is not significantly different from that of alloys in region B-I. This is not the case for region B-III, however, where a steep drop-off is observed in BCC% retention alongside further decreases in B as C_p increases. This trend is primarily driven by an increase in Ti concentration, which significantly weakens the mechanical strength and structural stability of the BCC phase. But since Ti possesses the highest pure metal heat capacity, increases in the concentration of Ti is likely necessary for maximal improvements in C_p . This is exactly what is seen in region B-IV of the Pareto front, where the Ti concentration has been maximized (by reducing Nb) to obtain the highest possible heat capacity and the greatest amount of ductility, albeit at the cost of BCC retention and stability during dynamic mechanical loading. Taken together, these results show that, compositionally, Mo is a significant source of strength and stiffness in this alloy, however, Ta and Ti are the most important when it comes to the tuning of material properties. Changes in Nb concentration seem to have little effect on properties and serve mostly to facilitate changes in the other elements.

We now move to the dynamic mechanical properties of Pareto optimal compositions. Figure 5 illustrates the time evolution of the shock compression for an alloy composition representative of each region (*i.e.*, B-I–B-IV) in the optimization B Pareto front. To help understand the response, hydrostatic pressure plots in are drawn in Fig. 5e, in alignment with each snapshot to show the position and direction of the shock wave front. In these plots, a positive stress indicates compression. Left-most column of Fig. 5 illustrates the initial compressive response of four alloys, slightly different elapsed times are taken due to variable wave speeds (alloy specific Hugoniot) but the position of the shock wave is aligned. The most deformation is seen immediately in Fig. 5c and d for alloys with compositions corresponding to regions B-III and B-IV of the Pareto front. The deformation in these structures is most likely due to extended shear banding³⁶ and the associated loss of atomic structural order (shown by blue atoms). Conversely, alloys with compositions in regions B-I and B-II of the Pareto front shown in Fig. 5a and b exhibit a temporary BCC→FCC (face-centered cubic) transition (shown by pink atoms) and the formation of small localized mirrored domains apparent to twin structures seen commonly in BCC structures.³⁷ In the second snapshot, a tensile rarefaction wave travels right-to-left after the compressive wave reflects at the right surface. For alloys with compositions belonging to regions B-I and B-II of the Pareto front, the FCC structure formed during initial compression has begun to revert back to BCC, however, the localized mirrored domains that were observed are still present and have grown in size. For the alloys belonging to regions B-III and B-IV of the Pareto front, the total deformation is far more significant and consists almost entirely of the disordering and amorphization of the lattice and the formation of small void regions. For the final snapshot, two pictures of the simulation cell are drawn, one with all atoms depicted, and one with all BCC atoms removed to better show the irreversibly deformed structure inside the simulation cell. At this final time snapshot, some ringing of the cell is present, but pressure profiles are appreciably close to zero to justify the BCC retention measurement. The difference in the remaining deformation within each structure is apparent. For alloys with compositions in region B-I of the Pareto front, almost all of the deformation observed in the earlier snapshots has reverted back into the pristine BCC phase, with only dispersed pockets of point defects remaining. For alloys with compositions in region B-II of the Pareto front, the residual disordering and defected structure is greater in volume, but a significant portion of the structure has returned to BCC. Finally, for alloys with compositions in regions B-III and B-IV of the Pareto front, on the other hand, the amount of totally disordered and amorphized structure has continued to grow even further relative to the prior snapshots and now accounts more than half of the total simulation cell in both cases. These results illustrate that the various regions composing the Pareto front in the optimization B spans very different deformation mechanisms and that as such these deformation mechanisms can be optimized as well. Combining these observations with the ones made in Fig. 4, we now have a deeper understanding of the structure-property relationships in the compositional space.

Interpretation of Alloy Performance Through the Electronic Density of States

Based on the results above, we performed electronic structure calculations of the density of states (DOS) using DFT to get a better understanding behind the compositional dependencies seen in the mechanical strength and stiffness. The electronic DOS contains valuable information from the relative position and filling of atomic orbitals in the valence and conduction bands of




Figure 5. Evolution of deformation during dynamic shock loading. Molecular dynamics shock columns demonstrating the deformation and restoration of the BCC lattice as a shockwave moves through (left-to-right) at 800 m/s. **a-d** Each row depicts a single simulation at different snapshots. Each composition is representative of typical alloy compositions belonging to each of the Pareto regions highlighted in Fig. 4 (Designated according to the same a-d naming convention). The three snapshots (in order from left-to-right) depict the time evolution of simulation cell, and show evolution of deformed structure during the transient shock event. Snapshots at t_1 and t_2 show the initial deformation caused by the compressive shockwave front, and initial deformation evolution as the lattice is decompressing. Snapshot t_f shows the deformation remaining when length of the shocked simulation cell has returned to its initial length from the start of the simulation, and is shown first with all atoms visible, and then again with all (grey colored) BCC atoms removed. **e** Hydrostatic pressure plots, where the x axis is aligned with the x position (visually) of simulation cells, indicate where the front shock wave is located within the cell and in which direction it is traveling. Note how the pressures in the third snapshot t_f fall appreciably close to zero just as the simulation cell is returning to its un-strained length, indicating that shockwave is momentarily absent from the cell.




Figure 6. Density of states for five representative alloy compositions. Electronic DOS for the each of the four **a-d** regions of the optimization B Pareto front in Figs. 4 and 5 (*i.e.*, regions B-I through IV). The DOS in **e** corresponds to the equiatomic composition as a reference point. Formation energies, calculated via DFT, show correlation between the structural and mechanical performance with the position (*i.e.*, the occupation level, n_F) of the Fermi energy (E_F) within the pseudo-gap region that forms due to the splitting of bonding and anti-bonding states that accompanies bond hybridization.

metals. In Fig. 6, we plot the total electronic DOS of each of the selected alloy compositions for regions B-I-B-IV of the Pareto front in optimization B and compare it to the total DOS of the equiatomic composition (Fig. 6e). We observe that a prominent feature of the DOS in these MoNbTaTi systems is the presence of a sizable pseudo-gap in the electronic density near the Fermi level. First characterized by Mott,³⁸ pseudo-gaps such as those observed in this figure are not uncommon in alloys consisting of transition metals from the d -block, in which the localization of electrons in the valence d -band can cause the hybridization of bonds between atoms. This effect is particularly strong in alloys rich in Mo and W, which have strong effective nuclear charges brought on by relativistic effects that cause a contraction of atomic radius,^{39,40} increasing the level of localization in the d -shell. The pseudo-gap forms as a result of the separation of bonding and anti-bonding states around the Fermi level that form as a result of this hybridization.^{41–43} When present in transition metal alloys, this pseudo-gap is often associated with low heat of formation (*i.e.*, increased structural stability), stemming from the more covalent-like and stiff nature of the hybridized bonds.⁴⁴ The influence of the pseudo-gap on mechanical stability in these alloys comes down to the positioning of the Fermi level (E_F) within the gap. When E_F lies exactly at the bottom of the pseudo-gap's well, an ordered intermetallic structure is usually preferred. When E_F lies on the right side of the well, d -electrons are driven into anti-bonding states and structures typically become unstable and susceptible to phase change. Conversely, when E_F sits on the left side of the well, disordered states can remain favorable, however, structural stability typically decreases as E_F moves farther away from the well. In our previous investigation of MoNbTaTi alloy,¹⁷ we showed that this pseudo-gap was present in hand-picked compositions rich with Mo, which were also associated with high structural stability, elastic stiffness and mechanical strength. The gap was entirely absent in compositions with low concentrations of Mo which tended to exhibit weaker elastic and mechanical strengths, along with poorer thermodynamic stability.

In Fig. 6, we observe that a pseudo-gap is present in each of the representative compositions. This result makes sense considering the high Mo content for all cases. The smallest pseudo-gap found amongst all the compositions tested belonged to the equiatomic composition. We note that the Fermi level, E_F , shifts further away from the well of the pseudo-gap (and therefore the expected structural stability typically decreases) as we navigate from compositions lying on the Pareto front in region B-I (high strength, high stiffness) to region B-IV (low strength, high ductility). This trend is simply a result of there being not enough electrons to fully stabilize the structure by pushing E_F into the pseudo-gap region.⁴⁵ Of the elements in this alloy, Mo has the highest number of valence d electrons. As such, high concentrations of Mo drive stability by not just instigating the presence of the pseudo-gap, but also by providing the electrons needed to push E_F closer to the more stable region of the pseudo-gap well. Conversely, Ti has the fewest valence electrons. In this case, alloys richer in Ti will be thermodynamically less stable, even in presence of a pseudo-gap created by a high Mo content, because there are not enough valence electrons to generate more stable bonding. To further demonstrate this effect, we have calculated formation energies, E_{form} (see Methods for description) for each of the alloys in Fig. 6, finding a near exact agreement between formation energy and position of E_F inside the pseudo-gap well, where lower (*i.e.*, more stable) energies correspond to Fermi levels sitting closer to the center of the gap. Interestingly, the pseudo-gap phenomenon discussed here for transition metals, is the same phenomenon important for the formation of superconducting states in some materials.⁴⁶ In particular, it has been found in these materials, that when a pseudo-gap is formed and the electron occupation level at the Fermi energy decreases (*i.e.*, sits nearer to the center of the gap), the heat capacity also goes down,⁴⁷ aligning exactly with the optimization results presented here.

Previous investigations suggest supercell DFT methods are highly adequate for modeling the elastic and mechanical properties in the MoNbTaTi alloy system^{17,48,49}, and so even though experimental validation in this study is limited by the scarcity of data in literature, we are confident that both the DFT and molecular dynamics descriptions of the alloy properties here are accurate and suitable for the optimization objectives of this work. In Supplementary Note 4, we show that DFT, MD, and rule-of-mixtures (ROM) predictions are remarkably consistent with one-another when used to evaluate the bulk moduli and heat capacities of the four Pareto regions of optimization B. Particularly, the trends (as in the change in property value from one composition to the next) exhibited by the DFT and MD models are nearly identical. Assuming that the DFT predictions are reliable benchmarks for accuracy, this suggests that the MD model is more than capable of assessing compositional property relationships in the aims of optimization.

Discussion

The results presented above demonstrated the power of multi-objective Bayesian optimization for not only the discovery of alloy compositions with targeted properties, but also to reveal complex structure-property relationships in the compositional space. Specifically, beyond identifying novel alloy compositions with tailored thermomechanical properties, our approach enabled the targeted selection of compositions exhibiting specific deformation mechanisms. Results for optimization B have shown that this is a powerful tool for optimizing complex alloy compositions when there are conflicting, antagonistic objectives. Moreover, the presence of a pseudo-gap in the electronic density near the Fermi level, identified through electronic structure calculations, aligns with observed mechanical and thermodynamic properties in these alloys. This agreement supports our optimization approach, which effectively explores the relationship between composition, structure, and desired properties by




Figure 7. Hypervolume (HV) improvement of optimizations A and B. Each line corresponds to a single sub-optimization loop. The dashed line represents the approximate point in each optimization in which effective convergence of the HV was achieved.

considering both thermomechanical properties and deformation mechanisms as performance objectives.

What these results are not necessarily showing is the efficacy of this approach to sample and explore the compositional space. Bayesian optimization provides a user with efficiency of next sample selection, requiring fewer sample evaluations to reach a desired level of optimization when compared with more traditional Edisonian approaches based on trial and error or human intuition. A useful tool in evaluating both the efficacy and progress of an optimization algorithm is hypervolume (HV) improvement. The HV is a performance indicator that serves as a measure of the volume dominated by the points of the Pareto front with respect to a reference point. Effectively serving as a measurement of the quality of a Pareto front, many multi-objective Bayesian optimization algorithms, including the TS-EMO algorithm, are constructed around the goal of achieving optimization through improvement of the HV. For further details of the HV and the general Bayesian optimization framework, the reader is referred to the Methods section. The HV performance indicators of the individual sub-optimization loops of optimizations A and B are shown in Fig. 7. For optimization A, all three individual optimization loops achieve convergence or near convergence of the Pareto front in ~ 10 -30 sample evaluations (beyond the initial dataset). Rapid improvement in the HV is observed in the first handful of evaluations. This meteoric rise in the HV is a result of the cooperative relationship shared between α_L , B , and the BCC% retention properties. The search algorithm is able to quickly identify the Pareto region of the performance space, and because this region is small and localized it is able to spend nearly every evaluation step probing in its vicinity. Only small gains in the HV are achieved after the first ~ 30 evaluations indicating convergence of the Pareto front. In optimization B, the HV improved more slowly in comparison. All three optimization loops showed significant increases in the HV within the first 5-10 evaluations, however, effective convergence was not achieved until ~ 50 -60 evaluations. In this case, the slower convergence is a result of the antagonistic relationship collectively shared between the C_P , B , BCC% properties and the expanded range and complexity of the Pareto region. Thus, single evaluations typically resulted in much smaller per-step improvements in the HV and a slower overall convergence of the Pareto front.

To demonstrate the efficacy of the TS-EMO algorithm for composition optimization, we performed a separate series of smaller compositional searches in which the performance of the Bayesian TS-EMO algorithm is compared against two other




Figure 8. A comparison of the efficacy of different search methods. **a-d** Hypervolume (HV) comparison of the three-objective searches mirroring optimization A, **(a)** and respective performance spaces resulting from **(b)** TS-EMO, **(c)** LHS, and **(d)** grid-based sampling schemes. **e-h** HVs and performance spaces of the searches mirroring optimization B. Each search was allowed 20 sample evaluations. TS-EMO searches were initialized from two starting points that corresponded to the equiatomic composition and a point just off the equiatomic, while grid and LHS points were chosen prior to the beginning of the search optimization. Because the order of evaluation for the LHS and grid techniques is irrelevant, the hypervolume is drawn as a straight line.

more traditional search methods: a Latin Hypercube sampling (LHS)⁵⁰ scheme and a grid-based sampling scheme. The grid-based scheme consisted of sampling the compositional space according to the formula $A_x(BCD)_{1-x}$ where elements A, B, C, D are sequentially chosen from the four elements of our quaternary alloy. Each scheme was applied to search the design space for two separate three-objective optimizations, limited to only 20 sample evaluations, that mirrored optimizations A and B in this work (*i.e.* α_L -B-BCC% and C_P -B-BCC% objective groupings). Importantly, the TS-EMO guided searches were seeded using only two initial samples so that the cumulative knowledge available to the algorithm at start was very low. The optimization performance of the three search methods is compared in Fig. 8 **a** and **b**, where the HV's obtained for the full set of grid and LHS samples are plotted alongside the per-step HV improvement of the TS-EMO optimization. The HV's obtained by the grid and LHS searches are plotted as constant lines, corresponding to the maximum HV achieved through evaluation of the full set of chosen compositions. Constant lines are drawn for the non-Bayesian searches because the order of evaluation is irrelevant, as the evaluation of all samples (in the chosen set of compositions) will always lead to the same HV.

For both optimizations (A and B), the HV of the TS-EMO algorithm quickly surpasses the HVs of the grid and LHS sample sets. In fact, in the case of the α_L -based optimization (Fig. 8**a**), in which all three objectives shared a cooperative relationship, the TS-EMO algorithm only needed a single sample evaluation to not just surpass the HVs of the entire 20 samples of the grid and LHS sets, but to double them both. For the C_P -based optimization (Fig. 8**b**), in which objectives shared an antagonistic relationship, the TS-EMO algorithm, while slower, only needed 7 sample evaluations to surpass the other sampling schemes and by the end of the 20 sample optimization, the HV was effectively double that the other two methods. To visualize this efficacy, the performance spaces belonging to each of these searches are shown in Fig. 8 **b-d** and **f-h**, where the targeting optimization of the Pareto regions by the TS-EMO algorithm vastly outperforms the grid and LHS searches. For the grid and LHS searches, desirable regions of the performance space were occasionally sampled, however, these regions still typically were far from the true Pareto fronts discovered by the TS-EMO algorithm.

Moving forward, there are a few potential technical challenges that can hinder the applicability of Bayesian optimization models to more complex systems for wider adoption and effectiveness. For example, in optimization problems involving a high number of objectives, performance can easily become hindered if the objectives exhibit complex relationships, or if more objectives are included. This is the curse of dimensionality, requiring a prohibitively large number of sample evaluations to effectively sample the complex and widely ranging Pareto optimal regions that can arise in many-dimensional performance domains. On the other hand, the surrogate Gaussian process (GP) functions used to model each objective property are quite

robust with respect to the number of design variables^{20,51} and so GP informed optimization models are ideally suited for the many-component HEA and CCA materials optimization problem. GPs are sensitive to outliers in the measured function data, however, and are generally incapable of modeling discontinuous function behavior in the objective space.⁵² Thus, the potential for any kind of zeroth or first order thermodynamic phase change within the searchable design space can severely limit the effectiveness and applicability of these Bayesian optimization models. Implementing constraints and stability considerations such as alloying from grain growth stabilization^{53,54} or the formation of specific phases based on complex phase diagram^{25,26} is therefore very much an open question.

The simulated results presented here provide a strong foundation for applications that involve experimental campaigns where optimization processes could not only focus on the discovery of materials with specific, targeted properties, but also on the discovery of emerging physical and chemical (dynamic) mechanisms. However such experimental campaigns face data-related challenges as they are costly and time-consuming, often producing noisy and sparse datasets. Additionally, managing uncertainty in these experiments is another critical challenge. Overall, while challenges remain, multi-objective Bayesian optimization holds great promise for accelerating alloy discovery by efficiently navigating the complex design space and identifying promising candidates with desired properties and deformation mechanisms.

METHODS

Molecular Dynamics Simulations




Figure 9. Calculations of objective functions. Molecular dynamics simulations are used to evaluate the various objective functions used on our Bayesian optimization workflow. **a** Coefficient of thermal expansion (α_L) is calculated by evaluating volume change with respect to temperature. **b** Heat capacity (C_P) is calculated as the change in enthalpy with respect to temperature. **c** Bulk modulus (B) is calculated by evaluating the changes in stress as a function of volume change. **d** Retention of BCC phase is evaluated through shock loading simulations.

We calculated the properties of the alloy (B , BCC%, α_L , C_P) as a function of composition to be used in our multi-objective Bayesian optimization scheme via molecular dynamics simulations. Examples of these calculations are illustrated in Fig. 9.

We performed all simulations using the Large-scale Atomic/Molecular Massively Parallel Simulator (LAMMPS) code.⁵⁵ In order to describe the MoNbTaTi alloy as a function of composition, for the (static) thermodynamic properties (*i.e.*, B , α_L , and C_P), we used the moment tensor potential (MTP) developed by Zheng *et al.*,⁵⁶ while for the shock loading simulations and the BCC% metric, we used a spectral neighbor analysis potential (SNAP) that we developed and trained specifically for modeling the effects of high volumetric and uniaxial strains over large temperature ranges.^{17,57} Usage of the MTP potential for B , α_L , and C_P calculations and the SNAP potential for the BCC% shock calculations was based on a thorough analysis and comparison of each potential's performance in modeling those properties. This analysis is provided in Note 1 of the Supplementary Information. B , BCC%, α_L , and C_P were calculated for every alloy composition in the initial dataset and for every subsequent composition predicted during the optimization process. Notably, we chose to use the MTP potential for static properties because it outperformed the SNAP potential in terms of accuracy, consistency and stability. The SNAP potential was used for shock calculations because the MTP potential failed to produce any deformed structure during shock loading. This is likely a result of there being very little data related to large compression or uniaxial deformation in the MTP's training data.

Bulk modulus: Bulk moduli were computed for each composition through the linear relationship between stress and strain, with $\sigma_{ij} = \mathbb{C}_{ijkl} : \epsilon_{kl}$. These simulations consisted of deforming the simulation in six different strain-controlled deformation pathways for the derivation of the elastic stiffness tensor ($\mathbb{C}_{ijkl}$). The bulk modulus was computed according to the Voigt-Reuss-Hill approximation where $B = (\mathbb{C}_{11} + 2\mathbb{C}_{12})/3$. Calculations for the bulk moduli were performed using a cubic simulation box containing 3 456 atoms, representing a $12 \times 12 \times 12$ multiplication of base 2-atom BCC unit cell.

Coefficient of thermal expansion and heat capacity: The linear coefficient of thermal expansion (α_L) and heat capacity at constant pressure, (C_P) were calculated using isothermal-isobaric (NPT) ensembles over increasing temperature intervals. The coefficient of thermal expansion was calculated as the change in the average lattice constant as the temperature is increased from 300 K to 1200 K, such that

$$\alpha_L = \frac{1}{L_0} \frac{\Delta L}{\Delta T}, \quad (1)$$

where L_0 is 300 K lattice constant. For this calculation, the lattice constant (L) was determined at three temperature intervals (300 K, 750 K, and 1200 K), and fit to a linear function where the slope, $m = \Delta L/\Delta T$. Likewise, the heat capacity was calculated according to the change in enthalpy as temperature is increased from 0 K to 300 K such that

$$C_P = \frac{\Delta H}{\Delta T}, \quad (2)$$

where the change in enthalpy, ΔH , is a function of internal energy, E , pressure p , and volume V , with $\Delta H = E + pV$. For all thermal expansion and heat capacity calculations, the simulation box used to model each composition contained 8 192 atoms, representing a $16 \times 16 \times 16$ unit cell multiplication, and was equilibrated at 300 K for 20 ps.

Shock-induced deformation: For the BCC% retention objective we performed atomistic shock simulations using a $22 \times 22 \times 150$ supercell column containing 145 200 atoms, with an initially perfect BCC lattice structure. This supercell size was chosen to facilitate short simulation times during the optimization process, and was selected based on a convergence of the BCC% parameter during testing of different supercell sizes (described in Supplementary Note 1). The initial structure was equilibrated prior to shock loading, using an NPT ensemble for 1 ps at a temperature 300 K. After this initial relaxation, periodic boundary conditions were kept only in the $\langle 100 \rangle$ and $\langle 010 \rangle$ directions, while the $\langle 001 \rangle$ boundaries were subjected to a non-periodic condition. To simulate a transient shock loading event, we used the built-in LAMMPS wall/piston fix to rapidly compress the column with an infinitely massive piston along the $\langle 001 \rangle$ direction, with one face being pushed towards the other at a constant velocity of 800 m/s. A shock wave would then propagate ahead of the piston through the elongated length of the column while the continuously applied piston would force the supercell to compress behind the shock wave. Upon reaching the far end of the simulation box, the shock wave would reflect and begin moving backwards through the column, while the far wall would be pushed outwards, causing the entire column to decompress and begin to elongate in the $\langle 001 \rangle$ direction. During this expansion the dimensions of the supercell briefly return to the initial box dimensions, whereupon, the fraction of remaining BCC phase in the structure is calculated using a polyhedral template matching (PTM) counting scheme. The temporal evolution of pressure was also computed via spatial binning along the shock direction. Additionally, we performed separate shock calculations for the four Pareto region compositions shown in Fig. 5 using larger supercells than those used for the optimizations. These supercells were $50 \times 50 \times 130$ replications of the base unit cell, contained 650 000 atoms, and used specifically to improve visualization for Fig. 5.

Density Functional Theory Simulations

We calculated the DOS of selected alloy compositions using DFT calculations. We performed all of our DFT calculations using the Vienna Ab-initio Simulation Package (VASP).^{58,59} Thus, periodic wave functions were approximated using plane-wave

basis sets and the description of the ion–electron interaction was modeled according to the projector-augmented wave (PAW) method.^{60,61} Electron exchange–correlation was approximated according to the generalized gradient approximation under the Perdew, Burke, and Ernzerhof (PBE) parametrization.⁶² The plane-wave energy cutoff for the electronic wave functions was set to 400 eV.

Atomic structures were constructed by randomizing the positions of atomic species across the BCC lattice sites. In place of a single quasi-random simulation, precision was achieved by modeling the random solid solution structure with three separately randomized supercells for each selected composition. These supercells contained 128 atoms corresponding to a $4 \times 4 \times 4$ replication of the base 2-atom BCC unit cell. Initial structural relaxation was performed using a conjugate gradient algorithm to relax atomic positions and the size and shape of the supercell. Partial orbital occupancy during initial relaxation was handled via Gaussian smearing with a smearing width of 0.1 eV. Minimization of the electronic wavefunctions was achieved using a 10^{-6} eV threshold for the difference between iterations, while ionic convergence was achieved when the norms of all forces between atoms fell below $0.02 \text{ eV.}\text{\AA}^{-1}$. For these initial structural relaxations, a gamma centered k-point mesh of $4 \times 4 \times 4$ was found to be sufficient in minimizing differences in the total energy. Once structurally relaxed supercells were obtained, formation energies, E_{form} , were calculated for each of the selected compositions, according to

$$E_F = \frac{1}{N} E_{\text{DFT}}^{\text{supercell}} - \sum x_i \mu_i, \quad (3)$$

where $E_{\text{DFT}}^{\text{supercell}}$ is average total energy of the three modeled supercells, N is the number of atoms in a single supercell, x_i is the composition of species i , and μ_i is the chemical potential of species i . Chemical potentials were calculated as the energy/atom of the groundstate structure of each species. For Mo, Nb and Ta this was the 2-atom BCC unit cell, while for Ti the groundstate was modeled as a 2-atom HCP unit cell.

For the accurate calculation of the electronic DOS, a self-consistent static calculation was performed on the structurally relaxed supercells. With the relaxed structure and the charge density of the self-consistent calculation, a second calculation was performed using a denser k-point mesh (gamma-centered, $8 \times 8 \times 8$) and a tetrahedron smearing method with Blöchl corrections and a width of 0.2 eV.

Bayesian Optimization Framework

The optimization goal of this work was to simultaneously optimize multiple material properties (*i.e.*, either through maximization or minimization) as a function of alloy composition. Within the Bayesian optimization framework, each objective property is modeled using a cheap and quick-to-evaluate surrogate model, from which a theoretical performance space can be built and sampled to determine the next set of compositional design parameters to be evaluated. Thus, the effective implementation of a Bayesian optimization algorithm is to build a surrogate model on known data, use this model to predict new design samples, and then update the surrogate model with measurements of the new design samples, repeating iteratively until a desired level of optimization is achieved.

Thompson-sampling efficient multi-objective optimization (TS-EMO): In this work we employ the algorithm developed by Bradford et al.,²⁷ commonly known as Thompson Sampling Efficient Multi-objective Optimization. As is typical of many optimization algorithms, surrogate objective functions are modeled in TS-EMO as Gaussian Processes (GP), which have the benefit of being both easy-to-fit and cheap-to-evaluate. In TS-EMO algorithm, hyperparameter tuning during the Gaussian process training step is performed via maximum a posteriori estimate (MAP) because it outperforms other algorithms like maximum likelihood estimate for small data sets. Where TS-EMO deviates from other common optimization algorithms is in its method of sampling the theoretical performance space, where it employs a Thompson sampling scheme based on spectral sampling of the GP functions. The objective of the sampling is to find a Pareto front of the theoretical performance space constructed by the GPs, which is achieved by performing a separate multi-objective optimization using the Non-Dominated Sorting Genetic Algorithm II (NSGA-II).⁶³ Spectral sampling of the GP functions is ideal here because it yields a linear predictor for each objective, which allows for very fast and efficient evaluation of the NSGA-II algorithm. Once a theoretical Pareto front (*i.e.*, a prediction of the true Pareto front) is obtained, an acquisition function guides the selection of the ideal point, or set of points in the batch selection case, based on a HV improvement criterion. The HV is a performance indicator that serves as a measure of the volume dominated by the points of the Pareto front calculated with respect to a reference point, and the acquisition function requires that the selected point(s) would most improve the HV if evaluated and found to be true. Using this approach the TS-EMO algorithm is able to optimally balance the trade-offs between exploration of unknown areas in the performance space with exploitation of known areas that lie close to the Pareto front.

It should be noted that the absolute value of the HV not necessarily important. Rather, only the relative improvements observed per step, indicating either optimization progress or convergence, are of note. The HV is simply a measure of a volume formed between the points of the Pareto front and a static reference point selected at the start of the optimization to lie just beyond the worst performing points of the initial dataset. As such, the magnitude of the value depends only on the explicit

values of the calculated objective properties. In this case, the performance data (*i.e.*, the property values) was normalized (from 0 to 1) prior to start of optimization. For details of the how the HV is calculated in this work and in TS-EMO algorithm, the reader is referred to references.^{64–66}

Constraining the design space: Pareto optimization of the material performance properties was achieved by tuning the global alloy composition variables. As configurational entropy is key to thermodynamic and mechanical stability in HEAs,^{6,8} it was necessary to limit the searchable design space to a region that encapsulates higher entropy configurations and atomic variation. Thus, design constraints were placed on the allowed concentration bounds of each atomic species. For Mo and Nb concentrations, the bounds were allowed to range from 5 at.% at minimum to 50 at.% at maximum. For Ta, the same minimum bound of 5 at.% was used, however, the maximum was limited to 32 at.%. This limiting of Ta concentrations was necessary due to a spontaneous breakdown in the performance of the SNAP potential discovered for alloy compositions rich in Ta (this breakdown is discussed further in Notes 2 and 3 of the Supplementary Information). Lastly, because pure Ti is an HCP metal, and high concentrations of Ti could potentially destabilize the parent BCC phase,¹⁷ the allowed Ti concentration was limited to a minimum of 2 at.% and a maximum of 25 at.%.

While there are four individual components in the MoNbTaTi alloy, during the actual implementation of the optimization algorithms only the concentrations of the Mo, Nb, and Ta species were explicitly included as tunable design variables. Thus, the algorithm was tasked with optimizing only three variables, while the fourth implicit variable (*i.e.*, the Ti concentration) was found deterministically from the sum of the other three variables relative to unity. Defining the design space in such a way allows us also to avoid the complication of solving a problem constrained by linear equality. This is the case when the sum of all design variables must always be the same, or as in this work, equate to one or unity. Algorithms often struggle to identify valid design sets when confined to such a tight constraint, and so it regularly becomes useful to reformulate the problem so that it is instead defined by a set of more tractable inequality constraints (*i.e.*, identifying sets of values that are greater or lower than some lower and upper bounds, respectively). By leaving the Ti concentration deterministic, we were able to bound the problem to the allowed concentration bounds of the Ti species. In doing so, we constrained the problem to a parabolic function, $g(x) = (x_{\text{Ti}} - L)^2 - R(x_{\text{Ti}} - L)$, where x_{Ti} is the current Ti concentration, L is the lower bound allowed for Ti concentrations and R is the range allowed for Ti concentrations (*i.e.*, the upper bound minus the lower bound). In the TS-EMO implementation used in this work (and in most Bayesian optimization implementations in general) design constraints are considered passed when the function returns a negative value or failed when the function returns a positive value. Thus, in this work, when the concentration of Ti (determined relative to the other species concentrations as $x_{\text{Ti}} = 1 - (x_{\text{Mo}} + x_{\text{Nb}} + x_{\text{Ta}})$), lies outside of the allowed bounds, a positive value is returned by the parabolic constraint function, whereas when the Ti concentration fits within the allowed bounds, the parabolic function returns a negative and passing value.

Initial datasets: Outside of the creation of a small initial dataset, both optimizations were designed to be fully autonomous such that the TS-EMO composition predictions were fed directly into the LAMMPS calculations used to evaluate each objective property, the results of which were in-turn fed back into the TS-EMO algorithm where new predictions were made and the process continued iteratively. Because there is a small degree of inherent randomness in the optimization sampling process, we performed each optimization three times for a total of 6 separate optimizations to ensure convergence of the Pareto optimal regions. Each of these optimizations was run for 100 iterations, with 1 composition selected and evaluated per iteration. The initial database was constructed partially according to a grid sampling scheme based on the formula $A_x(\text{BCD})_{1-x}$, and partially according a space-filling Latin-hypercube sampling (LHS) scheme. The initial dataset for each optimization consisted of the same 35 compositions, 21 of which were sampled according to the grid scheme, with the remaining 14 coming from the LHS scheme. Altogether, 670 composition evaluations were performed in this work (335 total for optimizations A and B individually), with 1 940 distinct LAMMPS property evaluations.

DATA AVAILABILITY

The data that support the findings of this study are available from the corresponding authors upon reasonable request.

CODE AVAILABILITY

The codes used to calculate the results of this study are available from the corresponding author upon reasonable request.

ACKNOWLEDGEMENTS

The authors would like to thank Cosmin Safta from Sandia National Laboratories for their comments and review of this work prior to submission. The multi-objective Bayesian optimization and atomistic capabilities and computational resources are supported in part by the Center for Integrated Nanotechnologies, an Office of Science user facility operated for the U.S. Department of Energy. This article has been authored by an employee of National Technology & Engineering Solutions of

Sandia, LLC under Contract No. DE-NA0003525 with the U.S. Department of Energy (DOE). The employee owns all right, title, and interest in and to the article and is solely responsible for its contents. The United States Government retains and the publisher, by accepting the article for publication, acknowledges that the United States Government retains a non-exclusive, paid-up, irrevocable, world-wide license to publish or reproduce the published form of this article or allow others to do so, for United States Government purposes. The DOE will provide public access to these results of federally sponsored research in accordance with the DOE Public Access Plan <https://www.energy.gov/downloads/doe-public-access-plan>.

AUTHOR CONTRIBUTIONS

JS and RD proposed the research idea and designed the overall study. JS and SD implemented the TS-EMO algorithm. JS generate the DFT and atomistic data. MM and MW helped with the atomistic simulation and developed the SNAP used in this work. JS performed the numerical analysis of the results. RD provided overall guidance and financial resources for the project. All the authors analyzed the data, discussed the results, and contributed to writing and reviewing the manuscript.

COMPETING INTERESTS

The authors declare no competing interests.

ADDITIONAL INFORMATION

Supplementary Information is available for this paper.

References

1. Suri, A. K., Krishnamurthy, N. & Batra, I. Materials issues in fusion reactors. *J. Phys. Conf. Ser.* **208**, 012001, DOI: [10.1088/1742-6596/208/1/012001](https://doi.org/10.1088/1742-6596/208/1/012001) (2010).
2. Boyer, R. R., Cotton, J. D., Mohaghegh, M. & Schafrik, R. E. Materials considerations for aerospace applications. *MRS Bull.* **40**, 1055–1066, DOI: [10.1557/mrs.2015.278](https://doi.org/10.1557/mrs.2015.278) (2015).
3. de Leon, N. P. *et al.* Materials challenges and opportunities for quantum computing hardware. *Science* **372**, eabb2823, DOI: [10.1126/science.abb2823](https://doi.org/10.1126/science.abb2823) (2021).
4. Boyce, B. *et al.* Machine learning for materials science: Barriers to broader adoption. *Matter* **6**, 1320–1323, DOI: [10.1016/j.matt.2023.03.028](https://doi.org/10.1016/j.matt.2023.03.028) (2023).
5. Fowler, J. E., Kottwitz, M. A., Trask, N. & Dingreville, R. Beyond combinatorial materials science: The 100 prisoners problem. *Integr. Mater. Manuf. Innov.* 1–9, DOI: [10.1007/s40192-023-00330-6](https://doi.org/10.1007/s40192-023-00330-6) (2024).
6. Miracle, D. B. & Senkov, O. N. A critical review of high entropy alloys and related concepts. *Acta Mater.* **122**, 448–511, DOI: [10.1016/j.actamat.2016.08.081](https://doi.org/10.1016/j.actamat.2016.08.081) (2017).
7. Cantor, B. Multicomponent and high entropy alloys. *Entropy* **16**, 4749–4768, DOI: [10.3390/e16094749](https://doi.org/10.3390/e16094749) (2014).
8. Senkov, O., Miracle, D., Chaput, K. & Couzinié, J.-P. Development and exploration of refractory high entropy alloys – A review. *J. Mater. Res.* **33**, 3092–3128, DOI: [10.1557/jmr.2018.153](https://doi.org/10.1557/jmr.2018.153) (2018).
9. Luo, H., Li, Z., Mingers, A. M. & Raabe, D. Corrosion behavior of an equiatomic CoCrFeMnNi high-entropy alloy compared with 304 stainless steel in sulfuric acid solution. *Corros. Sci.* **134**, 131–139, DOI: [10.1016/j.corsci.2018.02.031](https://doi.org/10.1016/j.corsci.2018.02.031) (2018).
10. Suárez Ocaño, P. *et al.* Influence of cooling rate on the microstructure and room temperature mechanical properties in the refractory AlMo_{0.5}NbTa_{0.5}TiZr superalloy. *J. Alloy. Compd.* **949**, 169871, DOI: [10.1016/j.jallcom.2023.169871](https://doi.org/10.1016/j.jallcom.2023.169871) (2023).
11. Elmslie, T. A. *et al.* Magnetic properties of equiatomic CrMnFeCoNi. *Phys. Rev. B* **106**, 014418, DOI: [10.1103/PhysRevB.106.014418](https://doi.org/10.1103/PhysRevB.106.014418) (2022).
12. Couzinié, J.-P., Senkov, O. N., Miracle, D. B. & Dirras, G. Comprehensive data compilation on the mechanical properties of refractory high-entropy alloys. *Data Brief* **21**, 1622–1641, DOI: [10.1016/j.dib.2018.10.071](https://doi.org/10.1016/j.dib.2018.10.071) (2018).
13. Melia, M. A. *et al.* High-throughput additive manufacturing and characterization of refractory high entropy alloys. *Appl. Mater. Today* **19**, 100560, DOI: [10.1016/j.apmt.2020.100560](https://doi.org/10.1016/j.apmt.2020.100560) (2020).
14. Lee, C. *et al.* An experimentally driven high-throughput approach to design refractory high-entropy alloys. *Mater. Des.* **223**, 111259, DOI: [10.1016/j.matdes.2022.111259](https://doi.org/10.1016/j.matdes.2022.111259) (2022).

15. Mooraj, S. & Chen, W. A review on high-throughput development of high-entropy alloys by combinatorial methods. *J. Mater. Inform.* **3**, 4, DOI: [10.20517/jmi.2022.41](https://doi.org/10.20517/jmi.2022.41) (2023).
16. Elmslie, T. A. *et al.* Tuning the magnetic properties of the CrMnFeCoNi cantor alloy. *Phys. Rev. B* **108**, 094437, DOI: [10.1103/PhysRevB.108.094437](https://doi.org/10.1103/PhysRevB.108.094437) (2023).
17. Startt, J., Kustas, A., Pegues, J., Yang, P. & Dingreville, R. Compositional effects on the mechanical and thermal properties of MoNbTaTi refractory complex concentrated alloys. *Mater. Des.* **213**, 110311, DOI: [10.1016/j.matdes.2021.110311](https://doi.org/10.1016/j.matdes.2021.110311) (2022).
18. Zhang, Z., Li, M., Cavin, J., Flores, K. & Mishra, R. A fast and robust method for predicting the phase stability of refractory complex concentrated alloys using pairwise mixing enthalpy. *Acta Mater.* **241**, 118389, DOI: [10.1016/j.actamat.2022.118389](https://doi.org/10.1016/j.actamat.2022.118389) (2022).
19. Shin, D., Yamamoto, Y., Brady, M. P., Lee, S. & Haynes, J. A. Modern data analytics approach to predict creep of high-temperature alloys. *Acta Mater.* **168**, 321–330, DOI: [10.1016/j.actamat.2019.02.017](https://doi.org/10.1016/j.actamat.2019.02.017) (2019).
20. Frazier, P. I. A tutorial on Bayesian optimization. Preprint at: <https://doi.org/10.48550/arXiv.1807.02811>, DOI: [10.48550/arXiv.1807.02811](https://doi.org/10.48550/arXiv.1807.02811) (2018).
21. Arróyave, R. *et al.* A perspective on Bayesian methods applied to materials discovery and design. *MRS Commun.* **12**, 1037–1049, DOI: [10.1557/s43579-022-00288-0](https://doi.org/10.1557/s43579-022-00288-0) (2022).
22. Giles, S. A., Sengupta, D., Broderick, S. R. & Rajan, K. Machine-learning-based intelligent framework for discovering refractory high-entropy alloys with improved high-temperature yield strength. *npj Comput. Mater.* **8**, 235, DOI: [10.1038/s41524-022-00926-0](https://doi.org/10.1038/s41524-022-00926-0) (2022).
23. Maresca, F. & Curtin, W. A. Mechanistic origin of high strength in refractory BCC high entropy alloys up to 1900K. *Acta Mater.* **182**, 235–249, DOI: [10.1016/j.actamat.2019.10.015](https://doi.org/10.1016/j.actamat.2019.10.015) (2020).
24. Pedersen, J. K. *et al.* Bayesian optimization of high-entropy alloy compositions for electrocatalytic oxygen reduction. *Angew. Chem.* **133**, 24346–24354, DOI: [10.1002/ange.202108116](https://doi.org/10.1002/ange.202108116) (2021).
25. Khatamsaz, D. *et al.* Multi-objective materials Bayesian optimization with active learning of design constraints: Design of ductile refractory multi-principal-element alloys. *Acta Mater.* **236**, 118133, DOI: [10.1016/j.actamat.2022.118133](https://doi.org/10.1016/j.actamat.2022.118133) (2022).
26. Khatamsaz, D. *et al.* Bayesian optimization with active learning of design constraints using an entropy-based approach. *npj Comput. Mater.* **9**, 49, DOI: [10.1038/s41524-023-01006-7](https://doi.org/10.1038/s41524-023-01006-7) (2023).
27. Bradford, E., Schweidtmann, A. M. & Lapkin, A. Efficient multiobjective optimization employing Gaussian processes, spectral sampling and a genetic algorithm. *J. Glob. Optim.* **71**, 407–438, DOI: [10.1007/s10898-018-0609-2](https://doi.org/10.1007/s10898-018-0609-2) (2018).
28. Baruffi, C., Maresca, F. & Curtin, W. A. Screw vs. edge dislocation strengthening in body-centered-cubic high entropy alloys and implications for guided alloy design. *MRS Commun.* **12**, 1111–1118, DOI: [10.1557/s43579-022-00278-2](https://doi.org/10.1557/s43579-022-00278-2) (2022).
29. Mishra, A., Kunka, C., Echeverria, M. J., Dingreville, R. & Dongare, A. M. Fingerprinting shock-induced deformations via diffraction. *Sci. Rep.* **11**, 9872, DOI: [10.1038/s41598-021-88908-y](https://doi.org/10.1038/s41598-021-88908-y) (2021).
30. Zhao, S. *et al.* Amorphization in extreme deformation of the CrMnFeCoNi high-entropy alloy. *Sci. Adv.* **7**, eabb3108, DOI: [10.1126/sciadv.abb3108](https://doi.org/10.1126/sciadv.abb3108) (2021).
31. Xie, Z. *et al.* Role of local chemical fluctuations in the shock dynamics of medium entropy alloy CoCrNi. *Acta Mater.* **221**, 117380, DOI: [10.1016/j.actamat.2021.117380](https://doi.org/10.1016/j.actamat.2021.117380) (2021).
32. Jian, W.-R., Xie, Z., Xu, S., Yao, X. & Beyerlein, I. J. Shock-induced amorphization in medium entropy alloy CoCrNi. *Scripta Mater.* **209**, 114379, DOI: [10.1016/j.scriptamat.2021.114379](https://doi.org/10.1016/j.scriptamat.2021.114379) (2022).
33. Ji, W. & Wu, M. S. Nanoscale origin of the crystalline-to-amorphous phase transformation and damage tolerance of Cantor alloys at cryogenic temperatures. *Acta Mater.* **226**, 117639, DOI: [10.1016/j.actamat.2022.117639](https://doi.org/10.1016/j.actamat.2022.117639) (2022).
34. Liu, S., Feng, G., Xiao, L., Guan, Y. & Song, W. Shock-induced dynamic response in single and nanocrystalline high-entropy alloy FeNiCrCoCu. *Int. J. Mech. Sci.* **239**, 107859, DOI: [10.1016/j.ijmecsci.2022.107859](https://doi.org/10.1016/j.ijmecsci.2022.107859) (2023).
35. Chase, M. W. & (US), N. I. S. O. *NIST-JANAF Thermochemical Tables*, vol. 9 (American Chemical Society, Washington, DC, 1998).
36. Bourne, N., Millett, J. & Gray, G. On the shock compression of polycrystalline metals. *J. Mater. Sci.* **44**, 3319–3343, DOI: [10.1007/s10853-009-3394-y](https://doi.org/10.1007/s10853-009-3394-y) (2009).

37. Ma, K. & Dongare, A. M. Role of $\alpha \rightarrow \epsilon \rightarrow \alpha$ phase transformation on the spall behavior of iron at atomic scales. *J. Mater. Sci.* **57**, 12556–12571, DOI: [10.1007/s10853-022-07381-8](https://doi.org/10.1007/s10853-022-07381-8) (2022).
38. Mott, N. F. Metal-insulator transition. *Rev. Mod. Phys.* **40**, 677, DOI: [10.1103/RevModPhys.40.677](https://doi.org/10.1103/RevModPhys.40.677) (1968).
39. Desclaux, J. P. & Kim, Y.-K. Relativistic effects in outer shells of heavy atoms. *J. Phys. B: At. Mol. Phys.* **8**, 1177, DOI: [10.1088/0022-3700/8/8/007](https://doi.org/10.1088/0022-3700/8/8/007) (1975).
40. Lassner, E. & Schubert, W.-D. *The Element Tungsten*, 1–59 (Springer US, Boston, MA, 1999).
41. Gelatt Jr, C., Williams, A. & Moruzzi, V. Theory of bonding of transition metals to nontransition metals. *Phys. Rev. B* **27**, 2005, DOI: [10.1103/PhysRevB.27.2005](https://doi.org/10.1103/PhysRevB.27.2005) (1983).
42. Pasturel, A., Colinet, C. & Hicter, P. Strong chemical interactions in disordered alloys. *Phys. B+C* **132**, 177–180, DOI: [10.1016/0378-4363\(85\)90062-2](https://doi.org/10.1016/0378-4363(85)90062-2) (1985).
43. Gourdon, O., Gout, D. & Miller, G. Intermetallic compounds, electronic states of. In Bassani, F., Liedl, G. L. & Wyder, P. (eds.) *Encyclopedia of Condensed Matter Physics*, 409–422, DOI: [10.1016/B0-12-369401-9/01117-7](https://doi.org/10.1016/B0-12-369401-9/01117-7) (Elsevier, Oxford, UK, 2005).
44. Lin, W., Xu, J.-H. & Freeman, A. Electronic structure, cohesive properties, and phase stability of Ni_3V , Co_3V , and Fe_3V . *Phys. Rev. B* **45**, 10863, DOI: [10.1103/PhysRevB.45.10863](https://doi.org/10.1103/PhysRevB.45.10863) (1992).
45. Ravindran, P. & Asokamani, R. Correlation between electronic structure, mechanical properties and phase stability in intermetallic compounds. *Bull. Mater. Sci.* **20**, 613–622, DOI: <https://doi.org/10.1007/BF02744780> (1997).
46. Shibauchi, T., Hanaguri, T. & Matsuda, Y. Exotic superconducting states in FeSe-based materials. *J. Phys. Soc. Jpn.* **89**, 102002, DOI: [10.7566/JPSJ.89.102002](https://doi.org/10.7566/JPSJ.89.102002) (2020).
47. Calegari, E., Magalhaes, S., Chaves, C. & Troper, A. pseudo-gap and the specific heat of high T_c superconductors. *Solid State Commun.* **158**, 20–24, DOI: [10.1016/j.ssc.2013.01.003](https://doi.org/10.1016/j.ssc.2013.01.003) (2013).
48. Coury, F. G. *et al.* Phase equilibria, mechanical properties and design of quaternary refractory high entropy alloys. *Mater. Des.* **155**, 244–256, DOI: [10.1016/j.matdes.2018.06.003](https://doi.org/10.1016/j.matdes.2018.06.003) (2018).
49. Coury, F. G., Kaufman, M. & Clarke, A. J. Solid-solution strengthening in refractory high entropy alloys. *Acta Mater.* **175**, 66–81, DOI: [10.1016/j.actamat.2019.06.006](https://doi.org/10.1016/j.actamat.2019.06.006) (2019).
50. McKay, M. D., Beckman, R. J. & Conover, W. J. A comparison of three methods for selecting values of input variables in the analysis of output from a computer code. *Technometrics* **42**, 55–61, DOI: [10.1080/00401706.2000.10485979](https://doi.org/10.1080/00401706.2000.10485979) (2000).
51. Moriconi, R., Deisenroth, M. P. & Sesh Kumar, K. High-dimensional bayesian optimization using low-dimensional feature spaces. *Mach. Learn.* **109**, 1925–1943, DOI: [10.1007/s10994-020-05899-z](https://doi.org/10.1007/s10994-020-05899-z) (2020).
52. Li, Z.-Z., Li, L. & Shao, Z. Robust gaussian process regression based on iterative trimming. *Astron. Comput.* **36**, 100483, DOI: [10.1016/j.ascom.2021.100483](https://doi.org/10.1016/j.ascom.2021.100483) (2021).
53. Murdoch, H. A. & Schuh, C. A. Stability of binary nanocrystalline alloys against grain growth and phase separation. *Acta Mater.* **61**, 2121–2132, DOI: [10.1016/j.actamat.2012.12.033](https://doi.org/10.1016/j.actamat.2012.12.033) (2013).
54. Monti, J. *et al.* Stability of immiscible nanocrystalline alloys in compositional and thermal fields. *Acta Mater.* **226**, 117620, DOI: [10.1016/j.actamat.2022.117620](https://doi.org/10.1016/j.actamat.2022.117620) (2022).
55. Thompson, A. P. *et al.* LAMMPS—a flexible simulation tool for particle-based materials modeling at the atomic, meso, and continuum scales. *Comput. Phys. Commun.* **271**, 108171, DOI: [10.1016/j.cpc.2021.108171](https://doi.org/10.1016/j.cpc.2021.108171) (2022).
56. Zheng, H. *et al.* Multi-scale investigation of short-range order and dislocation glide in MoNbTi and TaNbTi multi-principal element alloys. *npj Comput. Mater.* **9**, 89, DOI: [10.1038/s41524-023-01046-z](https://doi.org/10.1038/s41524-023-01046-z) (2023).
57. McCarthy, M. J., Startt, J., Dingreville, R., Thompson, A. P. & Wood, M. A. Atomic representations of local and global chemistry in complex alloys. Preprint at: <https://doi.org/10.48550/arXiv.2303.04311>, DOI: [10.48550/arXiv.2303.04311](https://doi.org/10.48550/arXiv.2303.04311) (2023).
58. Kresse, G. & Hafner, J. Ab initio molecular dynamics for liquid metals. *Phys. Rev. B* **47**, 558, DOI: [10.1103/PhysRevB.47.558](https://doi.org/10.1103/PhysRevB.47.558) (1993).
59. Kresse, G. & Hafner, J. Ab initio molecular-dynamics simulation of the liquid-metal–amorphous-semiconductor transition in germanium. *Phys. Rev. B* **49**, 14251, DOI: [10.1103/PhysRevB.49.14251](https://doi.org/10.1103/PhysRevB.49.14251) (1994).
60. Blöchl, P. E. Projector augmented-wave method. *Phys. Rev. B* **50**, 17953, DOI: [10.1103/PhysRevB.50.17953](https://doi.org/10.1103/PhysRevB.50.17953) (1994).

61. Kresse, G. & Furthmüller, J. Efficiency of ab-initio total energy calculations for metals and semiconductors using a plane-wave basis set. *Comput. Mater. Sci.* **6**, 15–50, DOI: [10.1016/0927-0256\(96\)00008-0](https://doi.org/10.1016/0927-0256(96)00008-0) (1996).
62. Perdew, J. P., Burke, K. & Ernzerhof, M. Generalized gradient approximation made simple. *Phys. Rev. Lett.* **77**, 3865, DOI: [10.1103/PhysRevLett.77.3865](https://doi.org/10.1103/PhysRevLett.77.3865) (1996).
63. Deb, K., Pratap, A., Agarwal, S. & Meyarivan, T. A fast and elitist multiobjective genetic algorithm: NSGA-II. *IEEE Trans. Evol. Comput.* **6**, 182–197, DOI: [10.1109/4235.996017](https://doi.org/10.1109/4235.996017) (2002).
64. Emmerich, M. T. M., Deutz, A. H. & Klinkenberg, J. W. Hypervolume-based expected improvement: Monotonicity properties and exact computation. In *2011 IEEE Congress of Evolutionary Computation (CEC)*, 2147–2154, DOI: [10.1109/CEC.2011.5949880](https://doi.org/10.1109/CEC.2011.5949880) (2011).
65. Emmerich, M., Yang, K., Deutz, A., Wang, H. & Fonseca, C. M. *A Multicriteria Generalization of Bayesian Global Optimization*, 229–242 (Springer International Publishing, Cham, Germany, 2016).
66. Yang, K., Emmerich, M., Deutz, A. & Fonseca, C. M. Computing 3-D expected hypervolume improvement and related integrals in asymptotically optimal time. In *International Conference on Evolutionary Multi-Criterion Optimization*, 685–700, DOI: [10.1007/978-3-319-54157-0_46](https://doi.org/10.1007/978-3-319-54157-0_46) (Springer, 2017).