



# Physics in the Machine: Integrating Physical Knowledge in Autonomous Phase-Mapping

A. Gilad Kusne<sup>1,2\*</sup>, Austin McDannald<sup>1</sup>, Brian DeCost<sup>1</sup>, Corey Oses<sup>3</sup>, Cormac Toher<sup>3</sup>, Stefano Curtarolo<sup>3</sup>, Apurva Mehta<sup>4</sup> and Ichiro Takeuchi<sup>2,5</sup>

<sup>1</sup>Materials Measurement Science Division, National Institute of Standards and Technology, Gaithersburg, MD, United States, <sup>2</sup>Materials Science and Engineering Department, University of Maryland, College Park, MD, United States, <sup>3</sup>Mechanical Engineering and Materials Science Department and Center for Autonomous Materials Design, Duke University, Durham, NC, United States, <sup>4</sup>Stanford Synchrotron Radiation Lightsource, SLAC National Accelerator Laboratory, Menlo Park, CA, United States, <sup>5</sup>Maryland Quantum Materials Center, University of Maryland, College Park, MD, United States

## OPEN ACCESS

### Edited by:

Tim Snow,  
Diamond Light Source,  
United Kingdom

### Reviewed by:

Marco Buongiorno Nardelli,  
University of North Texas,  
United States  
Carlo Barbieri,  
University of Surrey, United Kingdom

### \*Correspondence:

A. Gilad Kusne  
aaron.kusne@nist.gov

### Specialty section:

This article was submitted to  
Condensed Matter Physics,  
a section of the journal  
Frontiers in Physics

Received: 15 November 2021

Accepted: 21 January 2022

Published: 16 February 2022

### Citation:

Kusne AG, McDannald A, DeCost B, Oses C, Toher C, Curtarolo S, Mehta A and Takeuchi I (2022) Physics in the Machine: Integrating Physical Knowledge in Autonomous Phase-Mapping. *Front. Phys.* 10:815863. doi: 10.3389/fphy.2022.815863

Application of artificial intelligence (AI), and more specifically machine learning, to the physical sciences has expanded significantly over the past decades. In particular, science-informed AI, also known as scientific AI or inductive bias AI, has grown from a focus on data analysis to now controlling experiment design, simulation, execution and analysis in closed-loop autonomous systems. The CAMEO (closed-loop autonomous materials exploration and optimization) algorithm employs scientific AI to address two tasks: learning a material system's composition-structure relationship and identifying materials compositions with optimal functional properties. By integrating these, accelerated materials screening across compositional phase diagrams was demonstrated, resulting in the discovery of a best-in-class phase change memory material. Key to this success is the ability to guide subsequent measurements to maximize knowledge of the composition-structure relationship, or phase map. In this work we investigate the benefits of incorporating varying levels of prior physical knowledge into CAMEO's autonomous phase-mapping. This includes the use of *ab-initio* phase boundary data from the AFLOW repositories, which has been shown to optimize CAMEO's search when used as a prior.

**Keywords:** machine learning, phase mapping, autonomous physical science, scientific AI, phase diagram

## INTRODUCTION

Machine learning (ML) application into the physical sciences poses interesting challenges of data sparsity, high data collection cost, high data complexity, and learning intricate functional relationships. Regarding data cost and sparsity, obtaining new data involves performing very complex, resource-intensive, and time-consuming experiments in the lab or *in silico*. Performing a successful experiment requires hours to months of expert time using equipment often costing hundreds of thousands to millions of dollars (e.g., microdiffraction at synchrotron beamlines). Additionally, the expertise needed is measured in years past doctorate graduation. As a result, many physical science ML challenges must learn from a small number of observations. Furthermore, obtaining the target information such as the stoichiometric composition of a material with optimal properties, may require mapping the relationship between numerous input parameters and target variables; i.e., the relationship between elemental composition and functional properties. With each

new input parameter, the number of potential experiments grows exponentially. Consequently, the data obtained from costly experiments only sparsely represent a vast space of all possible experiments.

Confounding factors also include data complexity and the complexity of target relationships to be learned. Physical science data is often information rich. For instance, a Laue diffraction image from a material specimen contains information not only about crystal structures present in the sample, but also about their distribution, orientations, grain sizes, and crystallinity. Poor signal to noise and measurement setup-based signals, such as peaks due to the Cu K- $\beta$  spectral profile which may vary from instrument to instrument, can overwhelm features of interest. As a result, combining data from multiple instruments and studies can be highly involved. Furthermore, the relationships investigated with this data tend to be complex. This is particularly true of many technologically relevant materials; for example, the relationship between a ferroelectric material's microstructure and its piezoelectric response.

These challenges are often not shared by non-science application domains in which common ML methods arose, such as deep learning. For these domains, semi-uniform data and labels can be collected rapidly and cheaply. For instance, labels for text and object images are freely provided by internet service users seeking to prove that they are not bots [1]. No specialized expertise or equipment is needed, and data collection occurs in seconds. As a result, big data velocities and volumes are possible. The range of possible data for these domains is also bounded; for example, text images are bound by language and handwriting, car navigation is bound to roads, and chess moves are bound by the rules of the game. Typically, the goal is to optimize safely within these bounds, while scientific studies seek to explore edge cases.

Despite the additional challenges, science has a key advantage relative to common application domains—there are hundreds of years of literature containing theory and heuristics for guiding research. Scientific artificial intelligence (AI) focuses on encoding these rules (i.e., inductive bias) into AI frameworks to ensure that analysis results and predictions obey the scientific rules, and are therefore physically meaningful [2]. Restricting the solution space may offer an additional benefit of increasing data analysis speed. Probabilistic scientific AI incorporates uncertainty quantification and propagation into the analysis to better inform scientific decision making.

Scientific AI offers significant benefits for autonomous physical research systems [3], where AI controls experiment design, simulation, execution, and analysis. For these systems, scientific AI can ensure that prior physical knowledge informs the selection of subsequent experiments, and that each experiment is selected to obtain maximal information. While scientific knowledge can be encoded at multiple levels of the autonomous AI pipeline—from data representation through the performance measure used to update model parameters—much of the reported successes use off-the-shelf machine learning methods. This includes active learning [4] algorithms—machine learning algorithms dedicated to optimal experiment design, which are used to determine each subsequent experiment to be performed. Applications of off-the-shelf active learning algorithms include the use of genetic optimization for carbon nanotube process optimization [5], Gaussian process upper

confidence bounds to optimize molecular mixtures for photocatalysis [6], and estimate optimization for CO<sub>2</sub> electrocatalysis [7]. These successes of easily integrable, off-the-shelf active learning create opportunities and physical platforms where scientific AI may provide even greater research acceleration.

Recent work by Kusne and coworkers [8] demonstrates an autonomous physical research system for accelerating composition-phase-mapping and materials optimization, specifically the identification of optimal compositions that maximize some desired properties within a targeted search space. The autonomous system is driven by CAMEO (closed-loop autonomous materials exploration and optimization). This scientific AI algorithm was placed in control of the Stanford Synchrotron Radiation Lightsource high-throughput diffraction system, guiding each subsequent x-ray diffraction experiment, resulting in the discovery of a best-in-class phase change memory material. CAMEO was shown to accelerate materials optimization compared to standard methods by exploiting the materials composition-structure-property relationship to guide subsequent experiments. Toward this goal, CAMEO performs active phase-mapping—investigating subsequent compositions that provide maximal knowledge of the composition-structure relationship as represented by the composition-phase map. The structural phase map is fundamental to materials optimization as functional property extrema tend to occur within specific phase regions (e.g., magnetism and superconductivity) or along phase boundaries (e.g., martensitic transformation and morphotropic phase-boundary piezoelectrics). Knowledge of the phase map is used to guide materials optimization toward more promising regions of the search space.

Active phase-mapping can be thought of as an exploratory task to learn the composition-structure relationship. The composition space is segmented into regions based on which phases are present. To improve the performance of active phase-mapping, multiple levels of scientific knowledge can be incorporated, including density functional theory (DFT) data from the AFLOW.org repositories [9, 10]. This work investigates the impact of varying levels of incorporated physical knowledge on active phase-mapping performance. A full list of the algorithms studied, their varying levels of incorporated physical knowledge, and how the physical knowledge is encoded is provided in the Methods **Table 1**. Performance is explored for the benchmark ternary materials system of Fe-Ga-Pd [11].

## DISCUSSION

For this study, the level of scientific information in the active phase-mapping algorithm is varied by two factors—the first being the phase-mapping method. The structural phase-mapping method consists of 1) identifying the composition-phase map for samples with measured composition and x-ray diffraction patterns and then 2) extrapolating to samples without measured diffraction. Two phase-mapping methods are investigated. The first method uses go-to, off-the-shelf ML methods for clustering and classification: agglomerative hierarchical cluster analysis (HCA) with a cosine dissimilarity measure applied to the diffraction patterns [12] and a first-nearest neighbor algorithm for extrapolating phase region labels across the composition space. The alternative method uses the scientific AI

**TABLE 1** | Scientific AI physical knowledge and encoding method.

Algorithm	Physical knowledge	Encoding method
<b>Data Analysis</b>		
HCA CAMEO Phase-mapping [8]	Diffraction similarity identified by peak location rather than intensity. Phase regions are contiguous and phase boundaries are continuous	Use of Cosine dissimilarity measure [12] 1. If two or more sets of vertices share the same phase region label but are not connected by vertex neighbors, differing labels are assigned to the disconnected sets. 2. The Markov Random Field smoothness constraint [15]
	Materials of similar synthesis and processing parameters have similar properties Abundances of phases is non-negative X-ray diffraction intensity is non-negative Soft Gibbs Phase Rule—Upper bound limit on number of constituent phases Identified endmembers should be physically realizable	1. Markov Random Field smoothness constraint [15] 2. Harmonic Energy Minimization for label propagation [16] Karush–Kuhn–Tucker conditions [17] Karush–KuhnTucker conditions [17] Upper limit on number of endmember limits allowed in each phase region
Phase-mapping Prior	DFT phase map is predictive of bulk phase diagram. Structure is a good predictor of functional property and vice versa	Volume constraint on identified/predicted endmembers Bayesian prior through similarity kernel For more information see Refs [8, 13] M1c Phase Mapping: Phase mapping prior.
<b>Knowledge Propagation</b>		
1-NN	Samples of similar composition are likely to have similar phase.	As more samples are measured, the distance between samples in composition space gets smaller, so neighbors are more likely to have similar structure.
HEM	Phase regions are cohesive. Quantified likelihood for each sample belonging to each phase region due proximity in composition	Graph representation of composition space. Label propagation through graph. Labels uncertainty propagation.
<b>Active Learning</b>		
Sequence	None	
10 % Sampling	Samples chosen to be well distributed in composition space	Samples evenly distributed across composition space.
Uniform Random Sampling	Sampling uniformly will give general coverage of the composition space.	
Risk Minimization	Each sample quantified for its potential impact on improving total phase map performance. Targets phase boundaries.	Minimize total phase region misclassification probability for the entire phase map.

**TABLE 2** | Phase-mapping methods in order of performance (descending) at iteration 27.

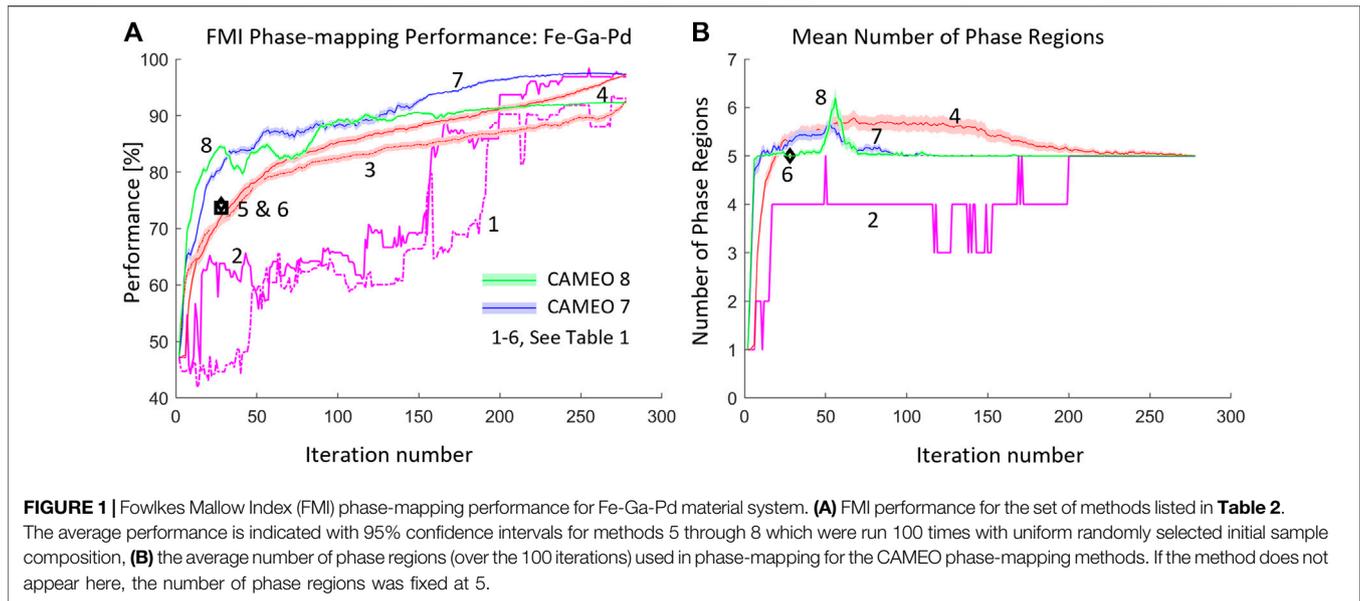
Algorithm index	Phase-mapping method	Prior	Active learning sampling method	Mean FMI performance for iteration 27 (%)
8, "CAMEO"	CAMEO Phase-mapping	Y	Risk Minimization	85
7, "CAMEO"	CAMEO Phase-mapping	N	Risk Minimization	80
6	CAMEO Phase-mapping	N	10%	74
5	HCA + 1NN	N	10%	74
4	CAMEO Phase-mapping	N	Random	72
3	HCA + 1NN	N	Random	71
2	CAMEO Phase-mapping	N	Sequence	64
1	HCA + 1NN	N	Sequence	45

HCA, hierarchical cluster analysis.

1NN, 1-Nearest Neighbor.

phase-mapping method of CAMEO. The CAMEO phase-mapping method employs a Bayesian graph-based algorithm to identify the probability of each composition sample belonging to each structural phase region. As a result, this method can generate a list of structural phase diagrams and their likelihoods. The method selects the most likely phase diagram based on the given data.

The optimal experiment design (OED) algorithm is the second factor varied, determining the sequence of samples to measure for diffraction data. Four methods are employed, as list in the column "Active Learning Sampling Method" in **Table 2**. The first method measures samples sequentially by their composition spread index [see Supplementary Figure 4(b) of Ref. [8]]. The next method selects samples randomly using a uniform distribution over



composition—a common exploratory active learning benchmark when the goal is gaining global knowledge of a search space. The third method selects each subsequent sample so that it minimizes total expected phase region misclassification error, here described as risk minimization [8]. This method was shown to target subsequent measurements along uncertain portions of the structural phase boundaries. The used risk minimization method requires a graph-based data representation and as such can only be combined with the graph-based CAMEO phase-mapping method. The sequential, random, and risk minimization methods are also compared to the performance of selecting 10% of the composition spread samples that provide good composition space coverage [see Supplementary Figure 4(a) of Ref. [8]]. The 10% coverage method is expected to provide good exploratory sampling and provide similar performance to the uniform random sampling as averaged over many runs.

As an additional modality for introducing prior physical knowledge, a Bayesian probabilistic prior over the phase map is implemented. The prior is derived from DFT calculations for the bulk Fe-Ga-Pd phase diagram as calculated by AFLOW [9, 10], with phase boundary data resolved by the AFLOW-CHULL [13] module (see Supplementary Figure 2 of Ref. [8]). The probabilistic prior is graph-based, defining the probability of materials belonging to the same phase region, and as such is demonstrated only in combination with the graph-based CAMEO phase-mapping method and the risk minimization OED method.

Autonomous phase-mapping performance is shown in **Figure 1A** using the modified Fowlkes-Mallow Index (FMI) performance measure [8], comparing the machine learning based phase-mapping results with expert labeled results. Here performance is averaged over 100 runs with the plot indicating the average performance with 95% confidence intervals (except for the 10% coverage OED method). Each autonomous phase-mapping method

is indexed and described in **Table 2**. The index number corresponds to a rank of performance at iteration 27, where 10% of the samples have been measured, allowing for comparison with the 10% sampling method. This is also the earliest iteration at which CAMEO Method 8 achieves an average performance of 85%.

In investigating the relative performance, it is interesting to note that the methods first group by OED method and then by phase-mapping method. For each OED method, the more physics-informed CAMEO phase-mapping method out-performs the off-the-shelf alternative. A complicating factor is that the off-the-shelf method is limited to phase-mapping with 5 structural phase regions, while the CAMEO phase map method allows the number of phase regions to vary and converge to an optimal. To ensure that the increase in performance is not due to an increase in the number of phase regions, i.e., model complexity, the average number of phase regions over the 100 runs is provided in **Figure 1B**.

OED performance also increases with greater prior physical knowledge. While sequential OED (Methods 1 and 2) simply contains information of sample location on the wafer, the use of the random and 10% sampling OED (Methods 3 through 6) assume that greater coverage of the composition space will provide more phase map knowledge. Finally, risk minimization (Methods 7 and 8) provides the best performance, building on the assumption that the most informative samples lie along phase boundaries.

Of particular interest is the fact that introducing prior information from AFLOW of the Fe-Ga-Pd bulk DFT phase diagram calculation (Method 8) achieves superior performance at lower iterations and then converges to a performance beneath those achieved by other methods including the CAMEO Method 7. Initially, when few diffraction patterns have been measured, the strong prior provides a correcting bias. However, as more data is obtained, the DFT-based bias pulls away from the correct answer for the thin film composition phase map.

For active phase-mapping, an increasing amount of physics information incorporated in the scientific ML provides better performance. While this improvement is demonstrated for a 2-dimensional composition space (3-simplex), it is expected that improvements will be more significant when searching higher dimensional spaces, as structural phase boundaries become exponentially sparser with increasing number of dimensions [13, 14]. Similarly, the search for optimal materials becomes increasingly difficult. As a result, the use of physics-informed active phase-mapping — through a combination of experiments and ab-initio calculations — is expected to become ever more important in guiding the search for novel, advanced materials.

## METHODS

### M1: Scientific AI

### M2 Statistics and Performance Metrics Confidence Interval

The 95% confidence interval was computed for the variable of interest over 100 experiments at the given iteration with:

$$CI_{95} = \left(\frac{\sigma}{\sqrt{n}}\right)F^{-1}(p, \nu) \quad (1)$$

Where  $F^{-1}$  is the inverse of the Student's t cumulative distribution function,  $\sigma$  is the standard deviation,  $n = 100$  is the number of experiments,  $p = \{2.5\%, 97.5\%\}$ , and  $\nu = 99$  is the degrees of freedom.

### Phase-Mapping Performance

Phase-mapping performance is evaluated by comparing phase region labels determined by experts with those estimated by CAMEO for the entire phase map (after the knowledge propagation step). To evaluate system performance, the Fowlkes-Mallows Index (FMI) is used, which compares two sets of cluster labels. The equations are presented below for the expert labels  $l \in L$  and the ML estimated labels  $\hat{l} \in \hat{L}$ , where the labels are enumerated  $L \rightarrow \mathbb{N}$  and  $\hat{L} \rightarrow \mathbb{N}$ .

If the number of phase regions is taken to be too large by either the user or the ML algorithm while the phase-mapping is correct, some phase regions will be segmented into sub-regions with the dominant phase boundaries preserved. For example, peak shifting can induce phase region segmentation<sup>44</sup>. To ensure that the performance measures ignore such sub-region segmentation, each estimated phase region is assigned to the expert labeled

phase region that shares the greatest number of samples. The number of phase regions is monitored to ensure that increases in model accuracy are not driven by increases in model complexity.

Fowlkes-Mallows Index:

$$FMI = TP / \sqrt{(TP + FP)(TP + FN)} \quad (2)$$

$$TP = \frac{1}{2} \sum_i \sum_j (l_i = l_j \& \hat{l}_i = \hat{l}_j) \quad (3)$$

$$FP = \frac{1}{2} \sum_i \sum_j (l_i \neq l_j \& \hat{l}_i = \hat{l}_j) \quad (4)$$

$$FN = \frac{1}{2} \sum_i \sum_j (l_i = l_j \& \hat{l}_i \neq \hat{l}_j) \quad (5)$$

## M3. Implementation

The methods were implemented in MATLAB\*. Built-in functions were used for agglomerative hierarchical cluster analysis and 1-nearest neighbors.

NIST Disclaimer: Certain commercial equipment, instruments, or materials are identified in this report in order to specify the experimental procedure adequately. Such identification is not intended to imply recommendation or endorsement by the National Institute of Standards and Technology, nor is it intended to imply that the materials or equipment identified are necessarily the best available for the purpose.

## DATA AVAILABILITY STATEMENT

The original contributions presented in the study are included in the article/Supplementary Material, further inquiries can be directed to the corresponding author.

## AUTHOR CONTRIBUTIONS

AK performed the computations and analysis with input from IT, AM, AMD, and BD. SC, CO, and CT provided the density functional theory data for the computations. The authors wrote the text together.

## ACKNOWLEDGMENTS

The authors thank Xiomara Campilongo and Marco Esters for fruitful discussions.

## REFERENCES

- Von Ahn L, Blum M, Hopper NJ, Langford J. CAPTCHA: Using Hard AI Problems for Security. In: International conference on the theory and applications of cryptographic techniques; May 4-8, 2003; Warsaw, Poland. Heidelberg, Germany: Springer (2003). p. 294-311.
- DeCost BL, Hatrick-Simpers JR, Trautt Z, Kusne AG, Campo E, Green ML. Scientific AI in Materials Science: A Path to a Sustainable and Scalable Paradigm - IOPscience. *Machine Learn Sci Technol* (2020) 1: 033001. doi:10.1088/2632-2153/ab9a20
- Stach E, DeCost B, Kusne AG, Hatrick-Simpers J, Brown KA, Reyes KG, et al. Autonomous Experimentation Systems for Materials Development: A Community Perspective - ScienceDirect. *Matter* (2021) 4(9):2702-76. doi:10.1016/j.matt.2021.06.036
- Settles B. *Active Learning Literature Survey*, United States: University of Wisconsin-Madison Department of Computer Sciences (2010). p. 11.
- Nikolaev P, Hooper D, Webber F, Rao R, Decker K, Krein M, et al. Autonomy in Materials Research: A Case Study in Carbon Nanotube Growth. *Npj Comput Mater* (2016) 2:16031. doi:10.1038/npjcompumats.2016.31

6. Burger B, Maffettone PM, Gusev VV, Aitchison CM, Bai Y, Wang X, et al. A mobile Robotic Chemist. *Nature* (2020) 583:237–41. doi:10.1038/s41586-020-2442-2
7. Zhong M, Tran K, Min Y, Wang C, Wang Z, Dinh C-T, et al. Accelerated Discovery of CO<sub>2</sub> Electrocatalysts Using Active Machine Learning. *Nature* (2020) 581:178–83. doi:10.1038/s41586-020-2242-8
8. Kusne AG, Yu H, Wu C, Zhang H, Hatrick-Simpers J, DeCost B, et al. On-the-fly Closed-Loop Materials Discovery via Bayesian Active Learning. *Nat Commun* (2020) 11:5966. doi:10.1038/s41467-020-19597-w
9. Oses C, Toher C, Curtarolo S. Data-driven Design of Inorganic Materials with the Automatic Flow Framework for Materials Discovery. *MRS Bull* (2018) 43:670–5. doi:10.1557/mrs.2018.207
10. Yang K, Oses C, Curtarolo S. Modeling Off-Stoichiometry Materials with a High-Throughput Ab-Initio Approach. *Chem Mater* (2016) 28:6484–92. doi:10.1021/acs.chemmater.6b01449
11. Long CJ, Hatrick-Simpers J, Murakami M, Srivastava RC, Takeuchi I, Karen VL, et al. Rapid Structural Mapping of Ternary Metallic alloy Systems Using the Combinatorial Approach and Cluster Analysis. *Rev Sci Instrum* (2007) 78:072217. doi:10.1063/1.2755487
12. Iwasaki Y, Kusne AG, Takeuchi I. Comparison of Dissimilarity Measures for Cluster Analysis of X-ray Diffraction Data from Combinatorial Libraries. *npj Comput Mater* (2017) 3:1–9. doi:10.1038/s41524-017-0006-2
13. Oses C, Gossett E, Hicks D, Rose F, Mehl MJ, Perim E, et al. AFLOW-CHULL: Cloud-Oriented Platform for Autonomous Phase Stability Analysis. *J Chem Inf Model* (2018) 58:2477–90. doi:10.1021/acs.jcim.8b00393
14. Toher C, Oses C, Hicks D, Curtarolo S. Unavoidable Disorder and Entropy in Multi-Component Systems. *npj Comput Mater* (2019) 69. doi:10.1038/s41524-019-0206-z
15. Kusne AG, Keller D, Anderson A, Zaban A, Takeuchi I. High-throughput Determination of Structural Phase Diagram and Constituent Phases Using GRENDEL. *Nanotechnology* (2015) 26:444002. doi:10.1088/0957-4484/26/44/444002
16. Zhu X, Ghahramani Z, Lafferty J. Semi-supervised Learning Using Gaussian fields and Harmonic Functions. In: Proceedings of the Twentieth International Conference on International Conference on Machine Learning; August 21 - 24, 2003; Washington, DC USA. California, U.S: AAAI Press (2003). p. 912–9.
17. Kuhn HW, Tucker AW. Nonlinear Programming. In: *Traces and Emergence of Nonlinear Programming*. Heidelberg, Germany: Springer (2014). p. 247–58. doi:10.1007/978-3-0348-0439-4\_11

**Conflict of Interest:** The authors declare that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

**Publisher's Note:** All claims expressed in this article are solely those of the authors and do not necessarily represent those of their affiliated organizations, or those of the publisher, the editors and the reviewers. Any product that may be evaluated in this article, or claim that may be made by its manufacturer, is not guaranteed or endorsed by the publisher.

Copyright © 2022 Kusne, McDannald, DeCost, Oses, Toher, Curtarolo, Mehta and Takeuchi. This is an open-access article distributed under the terms of the Creative Commons Attribution License (CC BY). The use, distribution or reproduction in other forums is permitted, provided the original author(s) and the copyright owner(s) are credited and that the original publication in this journal is cited, in accordance with accepted academic practice. No use, distribution or reproduction is permitted which does not comply with these terms.