

Full length article

Accelerating active learning materials discovery with FAIR data and workflows: A case study for alloy melting temperatures

Mohnish Harwani ^a, Juan C. Verduzco ^b, Brian H. Lee ^b, Alejandro Strachan ^b*

^a Department of Computer Science, Purdue University, West Lafayette, IN, 47907, USA

^b School of Materials Engineering and Birk Nanotechnology Center, Purdue University, West Lafayette, IN, 47907, USA



ARTICLE INFO

Keywords:

Active Learning
FAIR data

ABSTRACT

Active learning (AL) is a powerful sequential optimization approach that has shown great promise in the discovery of new materials. However, a major challenge remains the acquisition of the initial data and the development of workflows to generate new data at each iteration. In this study, we demonstrate a significant speedup in an optimization task by reusing a published simulation workflow available for online simulations and its associated data repository, where the results of each workflow run are automatically stored. Both the workflow and its data follow FAIR (findable, accessible, interoperable, and reusable) principles using nanoHUB's infrastructure. The workflow employs molecular dynamics to calculate the melting temperature of multi-principal component alloys. We leveraged all prior data not only to develop an accurate machine learning model to start the sequential optimization but also to optimize the simulation parameters and accelerate convergence. Prior work showed that finding the alloy composition with the highest melting temperature required testing several alloy compositions, and establishing the melting temperature for each composition took, on average, multiple simulations. By developing a workflow that utilizes the FAIR data in the nanoHUB database, we reduced the number of simulations per composition to one and found the alloy with the lowest melting temperature testing only three compositions. This second optimization, therefore, shows a speedup of 10x as compared to models that do not access the FAIR databases.

1. Introduction

The principles of Findable, Accessible, Interoperable, and Reusable (FAIR) data [1] are critical to accelerate progress not only in materials science [2,3] but also across other physical sciences and engineering disciplines [4–6]. Significant progress has been achieved through the development of specialized data platforms [7–11] and additional nation-level facilities are called for [12]. While the economic benefits of FAIR data have been estimated [13], concrete examples showcasing its impact on research and discovery remain notably sparse. Furthermore, despite extensive efforts dedicated to refining optimization methods, the essential role of data sharing and the necessary infrastructure has often been overlooked. The vast majority of research data collected remains private, with the published subset typically embedded in formats like graphs, tables, or text that are not machine-actionable. These practices hamper innovation and lead to considerable duplication of efforts [2]. In this context, we illustrate how the reuse of a FAIR workflow

and its associated datasets can dramatically expedite materials discovery tasks, demonstrating a 10-fold increase in efficiency. Building upon prior datasets that adhere to FAIR principles provides a robust foundation for advanced computational techniques in materials discovery. The discovery process typically involves an optimization in a high-dimensional space, characterized by variables such as composition, processing conditions, and microstructure. Relying solely on traditional Edisonian approaches proves impractical and inefficient given the vast array of potential combinations. To address this challenge, a form of sequential optimization, active learning (AL), has emerged as a powerful tool. AL requires high-quality data to train models that are then used to select the most meaningful experiments to perform next. AL has been successfully applied in numerous materials science problems, significantly reducing the experiments and resources needed when compared to traditional methods [14–22]. The lack of FAIR data hinders the unimpeded application of AL to a wide range of problems.

* Corresponding author at: School of Materials Engineering and Birk Nanotechnology Center, Purdue University, West Lafayette, IN 47907, USA.
E-mail address: strachan@purdue.edu (A. Strachan).

Recent applications of AL paired with experimental setups in materials science have demonstrated its effectiveness in the discovery and optimization of new materials. For instance, Kusne et al. presented an autonomous materials research lab powered by a closed-loop autonomous system for materials exploration and optimization (CAMEO) [15]. The CAMEO system was integrated into the synchrotron beamline at the SLAC National Accelerator Laboratory and was tasked with finding the composition with the largest optical bandgap in the Ge-Sb-Te ternary, with a new optimal material composition, GST-467, discovered in only 19 iterations. Similarly, McDannald et al. developed an autonomous neutron diffraction explorer (ANDiE) [16], a system that allows measurements of materials magnetic transition behaviors to be guided by AL and probabilistic modeling. ANDiE showed a 5-fold reduction in the number of measurements required to find the transition temperature. Finally, Moon et al. demonstrated an application of AL to search novel perovskite oxides for oxygen evolution electrocatalysis [17]. Their workflow successfully identified an optimal composition that exhibits superior electrocatalytic performance with the lowest known intrinsic overpotential.

AL has also been used to enhance physics-based simulations, in particular in the development of interatomic potentials. Yoo et al. showed the importance of judiciously selecting atomic structures to augment the training set when developing interatomic potentials [18]. Similarly, Zhang et al. [19] developed the Deep Potential Generator (DP-GEN) that actively selects data points that improve model performance, optimizing for accuracy with a minimal amount of reference data. Kulichenko et al. [20] also used an AL approach that modifies the potential energy surface to favor regions with configurations for which the model is uncertain, such as those associated with rare events.

Finally, AL has found one of its most important applications with the screening of potential candidates to reduce the number of experiments or simulations required for an optimization. A study by Farache et al. [21] proposed an AL workflow to identify alloys with high melting temperatures from a predefined set of alloys using molecular dynamics (MD) simulations. Using a random forest (RF) model to predict the alloys' melting temperature and corresponding uncertainty, an AL process selects compositions to label (i.e. characterize via MD) through information acquisition functions. These alloys are then simulated, and the predictive model is retrained on this expanded set of 'known' alloys, now containing an additional alloy. The model then iterates by making an updated prediction over the 'unknown' alloys. Importantly, the AL loop launches an end-to-end, fully autonomous MD workflow to characterize the melting temperature of the desired alloy. This is implemented and published for online simulation [23] as a nanoHUB Sim2L [11]. Importantly, Sim2Ls automatically index and store all input-output pairs of the workflows into a database that also follows FAIR principles. The approach required testing only ~ 15 compositions to find the highest melting alloy out of 555 possible cases.

This paper extends prior work to identify alloys with high melting temperatures by (1) learning from the prior FAIR data to reduce the number of iterations required find a converged melting temperature value using MD, and (2) improving the AL cycle by starting from all available data from the prior work. We demonstrate that this approach not only significantly decreases the time needed to identify alloys with optimal melting temperatures ($\sim 10\times$ improvement over previous work [21]), but it also allows for data reuse in subsequent optimizations to find alloys with the lowest possible melting temperatures. The fundamental goal is to demonstrate the reusability of data across different material characteristic optimizations. A single tool can learn from different optimizations across one criterion to create a large queryable database, and use AL to navigate that database on different criteria, finding optimal alloys in fewer iterations. Throughout this work, we will reference the work by Farache et al. [21] as *Work 1*.

2. Methods

2.1. FAIR workflows and data: Sim2Ls and ResultsDB

Establishing robust data-sharing practices and employing advanced cyberinfrastructure to support them is essential to enhance research efficiency and foster greater innovation across disciplines. This work leverages nanoHUB's infrastructure for FAIR workflows, simulations, and data [11]. These resources enable best data practices with minimal effort with the goal of creating interconnected research environment where information is readily available and usable by researchers worldwide. Recent advances in the nanoHUB environment make scientific workflows and data FAIR. These technologies, Sim2Ls and ResultsDB, allow developers to share their research workflows and contribute to a database of open community-generated data.

Sim2Ls are end-to-end research workflows with formally declared inputs and outputs and are implemented as flexible Jupyter notebooks. Inputs and outputs have well-defined types and metadata. Research workflows can contain data processing steps, analysis, and/or computational simulations. Additionally, Sim2Ls simplify the process of setting up a containerized environment, guaranteeing workflow reproducibility across various platforms. Each tool results in an independent dataset with a digital object identifier (DOI) and metadata. Each record includes a serialized query identifier (SQUID), its creation date, and all inputs and outputs. This database can be accessed through a web interface or via a REST API. This approach not only enhances reliability but also fosters transparency in research methodologies. Results of successful workflow execution runs are automatically validated, indexed, and stored in the ResultsDB, an accessible database.

2.2. Molecular dynamics

This work builds on data from the *meltheas* Sim2L [23] described in detail in *Work 1*. This Sim2L calculates a melting temperature for multi-principal component alloys (MPCAs) containing Cr, Co, Cu, Fe, and Ni through a solid-liquid coexistence method using MD simulations [24]. This method seeks to find the temperature at which solid and liquid samples coexist in equilibrium. Atomic interactions are modeled using a many-body, embedded-atom-method (EAM) potential developed by Farkas and Caro [25]. While this potential was optimized for the elastic properties of near-equiatomic MPCAs, it was not designed for melting temperature calculations and tends to overestimate the melting temperature of ternary and quaternary alloys. As was done in *Work 1*, and to restrict the optimization to FCC MPCAs, the design space was limited to alloys where no element exceeds 50 at. %. We stress that our goal is the demonstrate the effect of FAIR data on AL optimization loops and not to recommend alloys for experimental testing. Systems are generated from the primitive FCC unit cell (lattice parameter = 3.56 Å) replicated $8 \times 8 \times 18$ to create a periodic simulation cell. The Z-direction is normal to the solid-liquid interface. Atom types are assigned randomly following the desired alloy composition.

To create an initial two-phase sample, each half of the system gets assigned a different temperature, with a solid region at temperature T_{sol} and the remainder region at temperature T_{liq} under isobaric conditions using two thermostats. These temperatures are designed to cause half of the sample to melt and after a short equilibration (10 ps long) the system is allowed to reach equilibrium using MD simulations in the isoenthalpic, isobaric (NPH) ensemble. The melting temperature corresponds to the system temperature if, after equilibration, the system includes both solid and liquid in steady-state. This approach is advantageous compared to other simulation methods, such as direct melting, as it avoids hysteretic effects [24,26]. A successful determination of melting temperature depends on the choice of T_{sol} and T_{liq} as well as the length of the simulation.

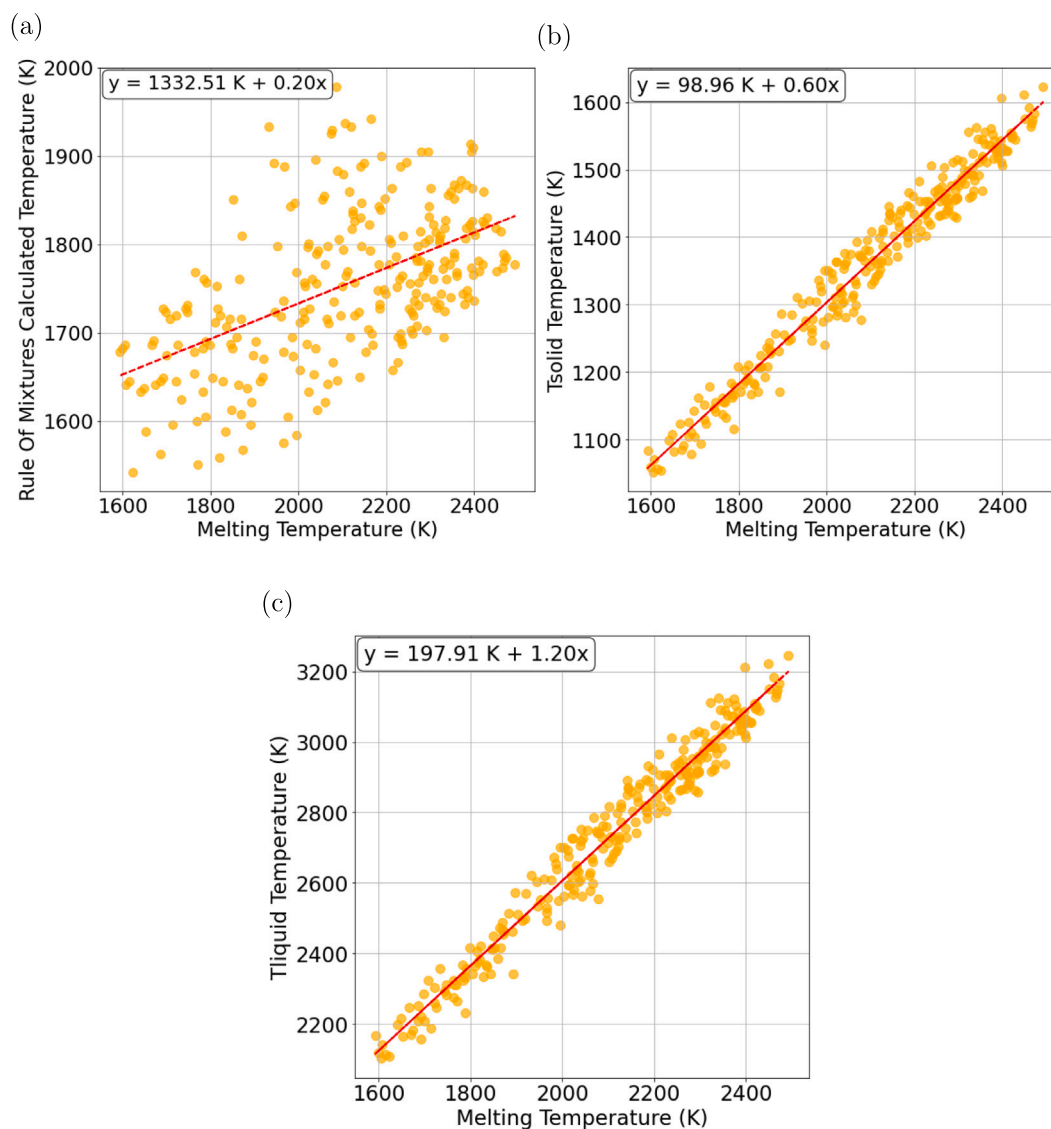


Fig. 1. Coexistence method melting temperature against (a) compositionally derived T_{ROM} ($R^2 = 0.2814$), (b) input temperature T_{sol} ($R^2 = 0.9884$), and (c) input temperature T_{liq} ($R^2 = 0.9884$).

2.3. Data collection and exploration

In *Work 1*, authors trained a RF model using a dataset of 39 alloy compositions to predict the melting temperature of multi-principal component alloys (MPCAs) containing Cr, Co, Cu, Fe, and Ni and efficiently explore the design space using AL. As an initial estimate for the T_{sol} and T_{liq} temperatures, they used a rule-of-mixtures calculation to estimate the melting temperature of the alloy in question and assigned T_{sol} 25% below this value and T_{liq} 50% above. These values were adjusted if liquid–solid coexistence was not achieved, the temperatures were increased or decreased by 5% depending on the final phase obtained, and the simulation was rerun until an acceptable result was obtained. In average, obtaining a converged value of the melting temperature required 4 simulations per composition. Across the different simulation times and acquisition functions explored in *Work 1*, a total of 265 different compositions were characterized, requiring 3039 simulations. Their design space consisted of 555 compositions.

We used the results automatically indexed in nanoHUB’s ResultsDB from the *MeltHEAs* Sim2L to improve on the simulations parameters and find the alloy composition with the minimum melting temperature in the fewest possible iterations. Since the nanoHUB Sim2L *MeltHEAs* is

open access, we took careful consideration to not include any compositions explored by the community after the original article publication date.

2.4. Input temperature modeling

Our first goal is the development of good estimators for T_{sol} and T_{liq} . Fig. 1(a) shows that, not surprisingly, the rule-of-mixtures estimate for the melting temperature used by Farache et al. (T_{ROM}) is not very accurate. Fig. 1(b–c) show T_{sol} and T_{liq} that resulted in converged MD results as a function of the MD-obtained melting temperature. Clearly, the data in the ResultsDB indicate that T_{sol} and T_{liq} values which lead to converged results can be obtained as a linear function of the alloy melting temperature.

Based on these observations, we hypothesized that accurate starting temperature values could be obtained from a predictive model of T_{melt} (alloy melting temperature) using simple linear models. Based on the approach outlined in *Work 1*, we used an 80/20 train/test split of their data to train a RF model to predict alloy melting temperatures. We kept their exact set of descriptors based on alloy composition (at. % of each element) and composition-derived features as inputs. This split resulted in a test set of 54 compositions. Our RF model included 100

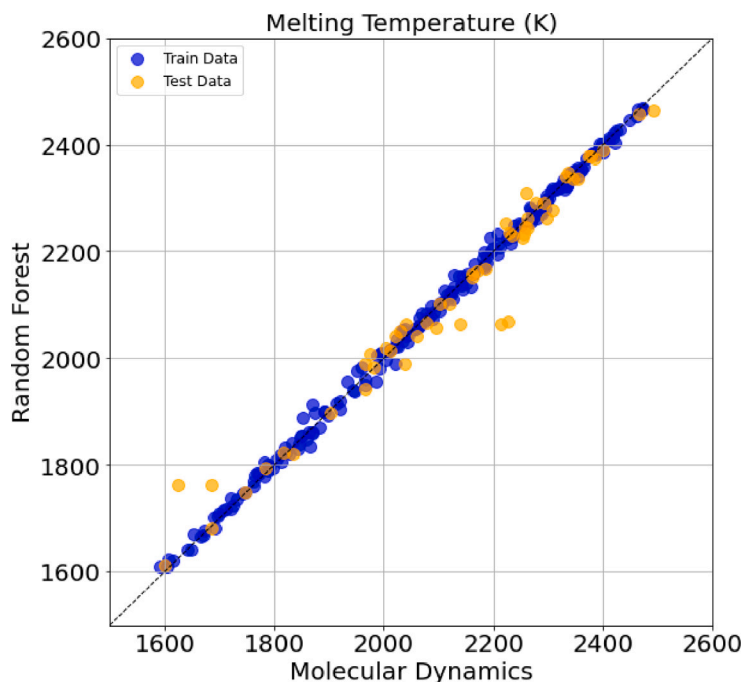


Fig. 2. Parity plot between Random Forest predictions of melting temperature and MD simulated melting temperature.

estimators, which were allowed to grow at to their maximum depth to reduce bias. Starting temperatures (T_{sol} and T_{liq}) were obtained from the RF (T_{melt}^{RF}) prediction using linear models derived from the training data. The predicted melting temperatures are shown in Fig. 2.

3. Results

3.1. Number of simulations to obtain a converged melting temperature for a given composition

We first evaluate our approach to estimate the T_{sol} and T_{liq} inputs to the MD simulations in a subset of previously explored alloys. Authors in *Work 1* estimated T_{sol} and T_{liq} from the RF prediction of the melting temperature (T_{melt}^{RF}) as $T_{liq} = 1.5T_{melt}^{RF}$ and $T_{sol} = 0.75T_{melt}^{RF}$. If the simulation did not result in coexistence and was either too cold or too hot, T_{liq} and T_{sol} were adjusted by increasing or decreasing T_{melt}^{RF} by 5%. Adjusting T_{liq} and T_{sol} often resulted in overcorrection, thus, in addition to the improved estimates, our approach uses a smaller adjustment of 2.5%, and, in the case of simulations that achieved coexistence (both fractions of solid and liquid alloy are between 35% and 65% of the system) but not a steady-state temperature measurement, the simulation time was adjusted by adding 50 ps (the default simulation time is 100 ps).

Using the sequestered test set from the RF training, 54 compositions, we analyzed the number of simulations required to achieve convergence and coexistence. The new model derived from the FAIR data in the ResultsDB clearly outperforms original approach. Fig. 3 shows that our model reduces the number of simulations per composition from an average of 4.4 to 1.3.

3.2. AL search of the alloy with the lowest melting temperature

To further demonstrate the efficiency improvements of our approach and available data, we used an AL workflow to find alloys with lowest melting temperature, as opposed to the search for the alloy with highest melting temperature explored in *Work 1*. The FAIR data in the ResultsDB enabled fine-tuning the simulations and the development of an accurate RF model to start the iterative optimization

resulting in significant decrease in the number of simulations needed to find an alloy with optimal melting temperature. To ensure a fair comparison, we modified our random forest to match the original implementation, using 350 estimators. Using the same acquisition functions modified to perform the reverse optimization (minimization), most acquisition functions were able to identify the alloy with the minimum melting temperature within 3 iterations using ~6 simulations total (approximately 2 simulations per alloy to determine its melting temperature).

Fig. 4 shows the AL iterations performed with the upper confidence bound (UCB) acquisition function. Figures S1–S3, describing the remaining acquisition functions, are included in the Supplemental information. We observed a high degree of similarity between alloys chosen by various acquisition functions. In addition to each acquisition function selecting 13 of the same alloys (out of 15 iterations), the order was remarkably similar. All alloys selected (across all acquisition functions) contained exactly 50% copper. Additionally, we noted that the acquisition functions tended to choose alloys very similar to each other in subsequent iterations, often only varying by 10% compared to the previous alloy.

Melting temperatures for the same alloy compositions were computed multiple times independently for each acquisition function, resulting in small deviations of approximately 10 K. Alloy compositions found to represent the minimum melting temperature were found within the budget for the acquisition functions, and are as follows: Cr 40%, Cu 50%, Ni 10% (1566 ± 11 K); Cr 40%, Cu 50%, Co 10% (1564 ± 9 K); and Cr 50%, Cu 50% (1535 ± 12 K). Uncertainty reported is the standard deviation between independent runs. Notably, the 50% Cu compositions we identified show melting temperatures consistent with the low values predicted for Cu-rich systems in *Work 1*, validating our AL results.

Comparing our melting temperature minimization problem to the original maximization in *Work 1* could be misleading if the two problems had different complexity. To address this, we also applied the original workflow (*Work 1*) to find the alloy with the lowest melting temperature using the original dataset of 40 alloys, stored in nanoHUB's ResultsDB, and preserved the same search space. The original task explored 15 compositions and used approximately 60 simulations (about 4 simulations per alloy) to identify the alloy with the highest melting

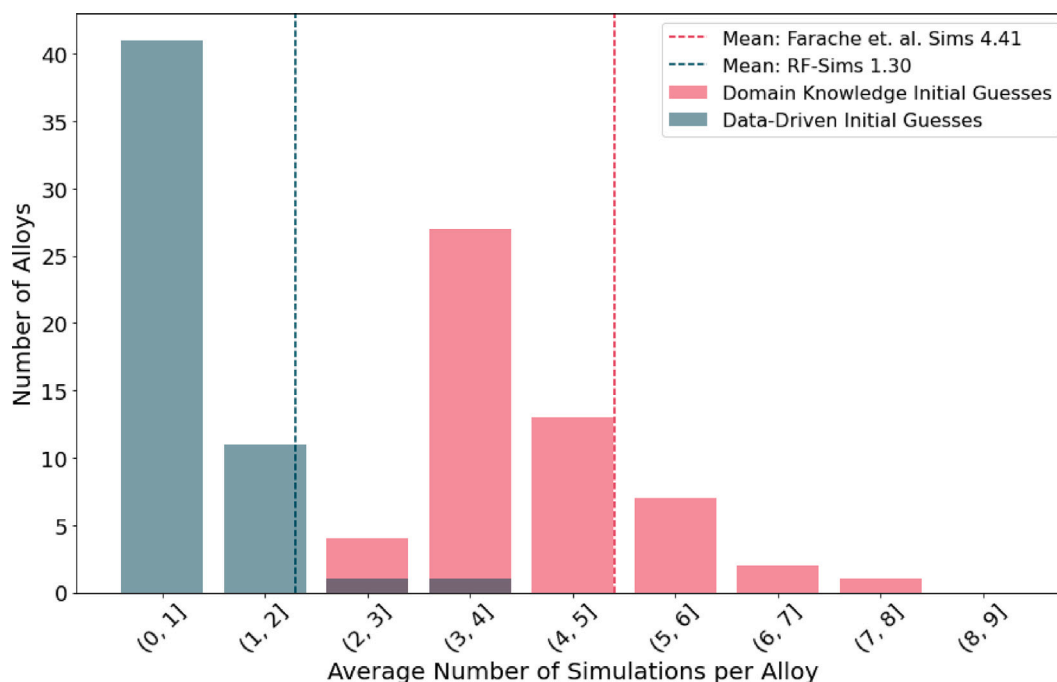


Fig. 3. Comparison of average number of simulations per alloy composition with and without a data-driven model for T_{liq} and T_{sol} temperatures.

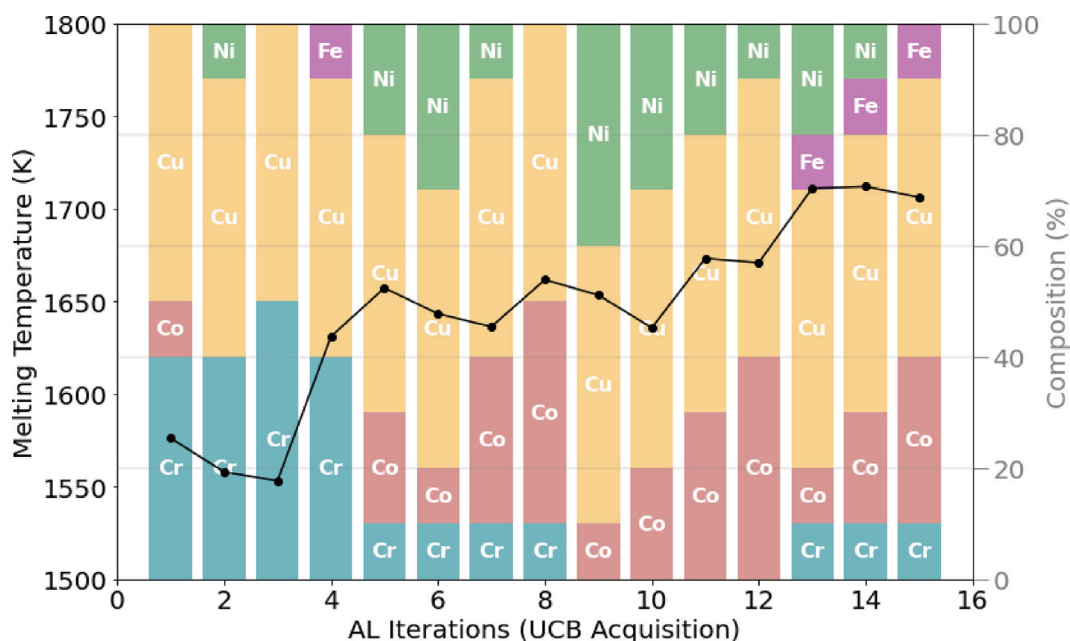


Fig. 4. Active learning run using the UCB activation function.

temperature from an unknown set of 555 compositions. The minimization task using the original training set and workflow required 16 compositions (65 simulations) using the UCB acquisition function. Using the MLI acquisition function, the alloy was not found within 40 compositions.

Overall, we demonstrate a significant increase in search efficiency in terms of number of compositions explored and number of simulations per composition. Across all acquisition functions tested, we demonstrate a 10x increase in performance compared to the results showcased in *Work 1*.

4. Conclusions

In this work, we demonstrated the transformative potential of FAIR (Findable, Accessible, Interoperable, and Reusable) data and workflows in computational materials science building on modern infrastructure. Building on a published tool for online MD simulations and prior results automatically captured we showed a 10-fold reduction in the resources required for the optimization of a material property as compared to our prior work, *Work 1* [21]. The latter required characterizing approximately 15 multi-principal component alloys to find the composition

with the highest melting temperature. Each composition required, in average, 4 simulations to find the melting temperature. We learned from the FAIR data collected during this initial effort to i) reduce the number of physics-based simulations required to determine the melting temperature of an alloy given a composition and ii) accelerate exploration and discovery of new alloys using active learning by starting with a more accurate model. Armed with a more accurate model and better starting parameters for the physics-based simulations we are able to find the alloy with the lowest melting temperature with only two simulations.

Our work provides a concrete example of how embracing FAIR data principles can accelerate innovation in materials science. We note that the general concepts and ideas apply well beyond the specific example shown here. While this work reuses a single FAIR workflow for a single quantity of interest, we envision future efforts combining data from multiple tools and various properties. Techniques such as transfer learning and multi-task learning could be used to learn from multiple sources and combine multiple tools.

CRedit authorship contribution statement

Mohnish Harwani: Writing – review & editing, Writing – original draft, Visualization, Software, Methodology, Formal analysis, Data curation. **Juan C. Verduzco:** Writing – review & editing, Writing – original draft, Supervision, Methodology. **Brian H. Lee:** Writing – review & editing, Writing – original draft, Supervision, Project administration, Methodology. **Alejandro Strachan:** Writing – review & editing, Writing – original draft, Validation, Supervision, Resources, Project administration, Methodology, Funding acquisition, Conceptualization.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgments

This effort was supported by the US National Science Foundation FAIROS program, award 2226418. This effort was also supported by the US National Science Foundation (DMREF-1922316). We acknowledge computational resources from nanoHUB and Purdue University through the Network for Computational Nanotechnology. We also thank Jorge Hernandez, Varun Vaidyanathan, and Kathleen O’Sullivan for their helpful discussion and efforts.

Appendix A. Supplementary data

Supplementary material related to this article can be found online at <https://doi.org/10.1016/j.commatsci.2024.113640>.

Data availability

The FAIR workflow described is available as a nanoHUB tool [27]. The simulation workflow to calculate melting temperature (meltheas) [23] is available as a Sim2L [11] in nanoHUB.

References

- [1] Mark D Wilkinson, Michel Dumontier, IJsbrand Jan Aalbersberg, Gabrielle Appleton, Myles Axton, Arie Baak, Niklas Blomberg, Jan-Willem Boiten, Luiz Bonino da Silva Santos, Philip E Bourne, et al., The FAIR guiding principles for scientific data management and stewardship, *Sci. Data* 3 (1) (2016) 1–9.
- [2] L Catherine Brinson, Laura M Bartolo, Ben Blaiszik, David Elbert, Ian Foster, Alejandro Strachan, Peter W Voorhees, Community action on FAIR data will fuel a revolution in materials research, *MRS Bull.* 49 (1) (2024) 12–16.
- [3] Matthias Scheffler, Martin Aeschlimann, Martin Albrecht, Tristan Bereau, Hans-Joachim Bungartz, Claudia Felser, Mark Greiner, Axel Groß, Christoph T Koch, Kurt Kremer, et al., FAIR data enabling new horizons for materials research, *Nature* 604 (7907) (2022) 635–642.
- [4] A Anil Sinaci, Francisco J Núñez Benjumea, Mert Gencturk, Malte-Levin Jauer, Thomas Deserno, Catherine Chronaki, Giorgio Cangioli, Carlos Caverro-Barca, Juan M Rodríguez-Pérez, Manuel M Pérez-Pérez, et al., From raw data to FAIR data: the fairification workflow for health research, *Methods Inf. Med.* 59 (S 01) (2020) e21–e32.
- [5] John Wise, Alexandra Grebe de Barron, Andrea Splendiani, Beeta Balali-Mood, Drashti Vasant, Eric Little, Gaspere Mellino, Ian Harrow, Ian Smith, Jan Taubert, et al., Implementation and relevance of FAIR data principles in biopharmaceutical R&D, *Drug Discov. Today* 24 (4) (2019) 933–938.
- [6] Toste Tanhua, Sylvie Pouliquen, Jessica Hausman, Kevin O’Brien, Pip Bricher, Caso De Bruin, Justin JH Buck, Eugene F Burger, Thierry Carval, Kenneth S Tacey, et al., Ocean FAIR data services, *Front. Mar. Sci.* 6 (2019) 440.
- [7] Anubhav Jain, Shyue Ping Ong, Geoffroy Hautier, Wei Chen, William Davidson Richards, Stephen Dacek, Shreyas Cholia, Dan Gunter, David Skinner, Gerbrand Ceder, et al., Commentary: The materials project: A materials genome approach to accelerating materials innovation, *APL Mater.* 1 (1) (2013).
- [8] James E Saal, Scott Kirklin, Muratahan Aykol, Bryce Meredig, Christopher Wolverton, Materials design and discovery with high-throughput density functional theory: the open quantum materials database (OQMD), *JOM* 65 (2013) 1501–1509.
- [9] Ben Blaiszik, Kyle Chard, Jim Pruyne, Rachana Ananthkrishnan, Steven Tuecke, Ian Foster, The materials data facility: data services to advance materials science research, *JOM* 68 (8) (2016) 2045–2052.
- [10] Claudia Draxl, Matthias Scheffler, NOMAD: The FAIR concept for big data-driven materials science, *MRS Bull.* 43 (9) (2018) 676–682.
- [11] Martin Hunt, Steven Clark, Daniel Mejia, Saaketh Desai, Alejandro Strachan, Sim2L: FAIR simulation workflows and data, *PLOS ONE* 17 (3) (2022) e0264492.
- [12] National Science and Technology Council (US), Materials Genome Initiative Strategic Plan, Executive Office of the President, 2021.
- [13] European Commission and Directorate-General for Research and Innovation, Cost-benefit analysis for FAIR research data – Cost of not having FAIR research data, Publications Office, 2018, <http://dx.doi.org/10.2777/02999>.
- [14] Burr Settles, Active Learning Literature Survey, University of Wisconsin-Madison Department of Computer Sciences, 2009.
- [15] A Gilad Kusne, Heshan Yu, Changming Wu, Huairuo Zhang, Jason Hattrick-Simpers, Brian DeCost, Suchismita Sarker, Corey Oses, Cormac Toher, Stefano Curtarolo, et al., On-the-fly closed-loop materials discovery via Bayesian active learning, *Nature Commun.* 11 (1) (2020) 5966.
- [16] Austin McDannald, Matthias Frontzek, Andrei T Savici, Mathieu Doucet, Efrain E Rodriguez, Kate Meuse, Jessica Opsahl-Ong, Daniel Samarov, Ichiro Takeuchi, William Ratcliff, et al., On-the-fly autonomous control of neutron diffraction via physics-informed Bayesian active learning, *Appl. Phys. Rev.* 9 (2) (2022).
- [17] Junseok Moon, Wiktor Beker, Marta Siek, Jiheon Kim, Hyeon Seok Lee, Taegwan Hyeon, Bartosz A Grzybowski, Active learning guides discovery of a champion four-metal perovskite oxide for oxygen evolution electrocatalysis, *Nature Mater.* 23 (1) (2024) 108–115.
- [18] Pilsun Yoo, Michael Sakano, Saaketh Desai, Md Mahbulul Islam, Peilin Liao, Alejandro Strachan, Neural network reactive force field for C, H, N, and O systems, *Npj Comput. Mater.* 7 (1) (2021) 9.
- [19] Linfeng Zhang, De-Ye Lin, Han Wang, Roberto Car, Weinan E, Active learning of uniformly accurate interatomic potentials for materials simulation, *Phys. Rev. Mater.* 3 (2) (2019) 023804.
- [20] Maksim Kulichenko, Kipton Barros, Nicholas Lubbers, Ying Wai Li, Richard Messerly, Sergei Tretiak, Justin S Smith, Benjamin Nebgen, Uncertainty-driven dynamics for active learning of interatomic potentials, *Nat. Comput. Sci.* 3 (3) (2023) 230–239.
- [21] David E Farache, Juan C Verduzco, Zachary D McClure, Saaketh Desai, Alejandro Strachan, Active learning and molecular dynamics simulations to find high melting temperature alloys, *Comput. Mater. Sci.* 209 (2022) 111386.
- [22] Juan C. Verduzco, Ernesto E. Marinero, Alejandro Strachan, An active learning approach for the design of doped LLZO ceramic garnets for battery applications, *Integrating Mater. Manuf. Innov.* 10 (2) (2021) 299–310.
- [23] Zachary D. McClure, Saaketh Desai, Alejandro Strachan, High entropy alloy melting point calculation, 2020, <http://dx.doi.org/10.21981/W5VD-T039>, URL: <https://nanohub.org/resources/meltheas>.
- [24] James R. Morris, C.Z. Wang, K.M. Ho, C.T. Chan, Melting line of aluminum from simulations of coexisting phases, *Phys. Rev. B* 49 (5) (1994) 3109.
- [25] Diana Farkas, Alfredo Caro, Model interatomic potentials and lattice strain in a high-entropy alloy, *J. Mater. Res.* 33 (19) (2018) 3218–3225.
- [26] James R. Morris, Xueyu Song, The melting lines of model systems calculated from coexistence simulations, *J. Chem. Phys.* 116 (21) (2002) 9352–9358.
- [27] Mohnish Harwani, Juan Carlos Verduzco, Alejandro Strachan, Active learning loop for MELTHEAS using FAIR principles, 2024, <http://dx.doi.org/10.21981/B5HH-DG46>, URL: <https://nanohub.org/resources/fairmeltheas>.