

RESEARCH ARTICLE

Benchmarking active learning strategies for materials optimization and discovery

Alex Wang ^{1,*}, Haotong Liang ¹, Austin McDannald ²,
Ichiro Takeuchi ¹ and Aaron Gilad Kusne ^{1,2}

¹Materials Science & Engineering Department, University of Maryland, College Park, MD, USA and ²Materials Measurement Science Division, National Institute of Standards and Technology, Gaithersburg, MD, USA

*Correspondence address. Department of Materials Science and Engineering, 2135 Chemical and Nuclear Engineering Building, 4418 Stadium Drive, University of Maryland, College Park, MD 20742-2115, USA. Tel: 301-405-5240; E-mail: aawang@umd.edu

ABSTRACT

Autonomous physical science is revolutionizing materials science. In these systems, machine learning (ML) controls experiment design, execution and analysis in a closed loop. Active learning, the ML field of optimal experiment design, selects each subsequent experiment to maximize knowledge toward the user goal. Autonomous system performance can be further improved with the implementation of scientific ML, also known as inductive bias-engineered artificial intelligence, which folds prior knowledge of physical laws (e.g. Gibbs phase rule) into the algorithm. As the number, diversity and uses for active learning strategies grow, there is an associated growing necessity for real-world reference datasets to benchmark strategies. We present a reference dataset and demonstrate its use to benchmark active learning strategies in the form of various acquisition functions. Active learning strategies are used to rapidly identify materials with optimal physical properties within a compositional phase diagram mapping a ternary materials system. The data are from an actual Fe-Co-Ni thin-film library and include previously acquired experimental data for materials compositions, X-ray diffraction patterns and two functional properties of magnetic coercivity and the Kerr rotation. Popular active learning methods along with a recent scientific active learning method are benchmarked for their materials optimization performance. Among the acquisition functions benchmarked, Expected Improvement demonstrated the best overall performance. We discuss the relationship between algorithm performance, materials search space complexity and the incorporation of prior knowledge, and we encourage benchmarking more and novel active learning schemes.

Key words: reference data; machine learning; phase map; materials optimization; benchmark; acquisition function.

INTRODUCTION

Technological advances are often dictated and driven by materials discovery. The need for ever-better materials spurs modern scientists to explore materials of greater and greater

complexity. For instance, interest in high-temperature superconductors has grown from the study of single-element materials to complex compounds such as Hg-Tl-Ba-Ca-Cu-O [1].

Submitted: 12 April 2022; Received (in revised form): 5 June 2022. Accepted: 26 June 2022

© The Author(s) 2022. Published by Oxford University Press.

This is an Open Access article distributed under the terms of the Creative Commons Attribution License (<https://creativecommons.org/licenses/by/4.0/>), which permits unrestricted reuse, distribution, and reproduction in any medium, provided the original work is properly cited.

Similarly, the number of elements used in the electronics industry increased from around 10–50 during the 1990s [2]. However, with each newly-added stoichiometric element or processing parameter, the number of possible materials to investigate grows exponentially. As a result, the traditional expert-driven Edisonian, one-by-one trial-and-error approach is rapidly becoming impractical.

In these Edisonian studies, materials scientists first select a target materials system (e.g. a ternary system A-B-C consisting of chemical elements A, B and C) to investigate, bounding the study to a composition and processing space. This in turn determines the experiment setup, such as the fabrication method, the material sources to use and appropriate materials processing equipment. Prior knowledge is then used to build a model, heuristic or intuition to predict desired material properties from materials synthesis parameters. The materials scientist then uses the predictive model to guide subsequent materials synthesis, characterization and analysis.

Combinatorial high-throughput (CHT) strategies were developed to enhance the rate and the efficiency of materials exploration [3]. CHT strategies allow for hundreds to thousands of materials from a target materials system (A-B-C) to be synthesized in parallel as a composition library. The library consisting of hundreds of different compositions $A_xB_yC_{1-x-y}$ (where, x and y are the compositional parameters varied on the library wafer layout with increment, for instance, of 0.01 (with $0 \leq 1-x-y \leq 1$)) is then loaded into a characterization system where each material is measured in rapid succession. However, for characterization methods of high cost or time, such as X-ray photoelectron spectroscopy or determination of the band gap, measuring all of the hundreds (or even sometimes thousands) of samples within a given library can be prohibitive. This challenge motivated the use of active learning, or the machine learning (ML) field of optimal experiment design (OED), to guide sequence of measurement experiments across the library [4]. Each measurement is selected to maximize knowledge toward a user-specific goal, e.g. identifying the compositional parameters x and y which gives the optimal physical property. An important associated goal is often to determine the composition-phase map across the entire ternary A-B-C. The use of active learning enables a more streamlined procedure for screening the materials phase space and provides the ability to do on-the-fly, adaptive and iterative learning and optimization with a minimal number of experiments [5]. Depending on the complexity of the composition-property landscape across the ternary, it is possible to arrive at the ‘correct answer’, i.e. the optimum composition after only a fraction of the entire library is measured. This aspect of active learning is particularly attractive when the search parameter space is extended to multiple dimensions beyond mapping of ternary phase diagrams.

The use of active learning in materials science gained popularity half a decade ago, with active learning driving recommendation engines to guide experimentalists in the lab [6]. Active learning is often combined with a ML surrogate model for when the underlying mechanistic model is unknown. Such active learning tools provided improved performance over traditional OED methods [7], as the predictive model (i.e. ‘response model’) is updated upon each iteration, ensuring always optimized decision-making [8]. More recently active learning has been integrated into autonomous materials research systems capable of performing experiment design, execution, and analysis in a closed-loop [9]. For example, autonomous systems have been used to optimize materials processing parameters to tune quantum dots optical behavior [10], or come up with best molecular mixtures for improved photovoltaics films [11], and identify the

new best-in-class phase change memory material—the first autonomous discovery of a best-in-class material [4]. The success of these autonomous systems depends on the active learning schemes employed.

Each active learning scheme pairs a predictive model, used to ‘forecast’ the properties of yet-to-be-measured materials with an acquisition function, which defines the utility of investigating each possible material. Materials compositions of maximal utility are then selected for subsequent studies. A probabilistic predictive model provides added advantage: these models output both an estimate and uncertainty for predictions, an example being the Gaussian Process (GP) [12] which is used here. Prediction estimate and uncertainty can then be combined in a Bayesian optimization (BO) algorithm to define utility [13]. When the task is materials optimization, successive BO acquisition functions strike a balance between an exploratory methods and exploitative methods. Exploratory methods seek global knowledge of the unknown target function, exemplified by choosing materials where the prediction model has maximum uncertainty. Exploitation methods search for optima of the target function, exemplified by choosing materials with predicted property optima. The combination allows an active learning scheme to avoid falling into local optima and provides greater speed and stability in the search for global optima [13]. The performance of each active learning scheme is dependent on the selection of the materials model, the acquisition function and the materials challenge and system being addressed.

As the number of active learning schemes increases, and the target materials challenges become more complex, the choice of active learning schemes becomes more important. The wrong selection can lead a researcher or autonomous system astray or greatly delay reaching the researcher’s goals, e.g. identifying a novel, optimal material. Materials datasets can be used to benchmark and select optimal active learning schemes. For the challenge of solid-state materials discovery and optimization, the relevant dataset would include data on materials synthesis such as composition, exhaustively varying these parameters and pairing them with the resulting structure and functional properties, as functional property is closely tied to structure. Using such a complete dataset, one can build a highly accurate composition-structure-property model, which can be used for the purpose of evaluating the efficacy of different active learning schemes through simulating closed-loop

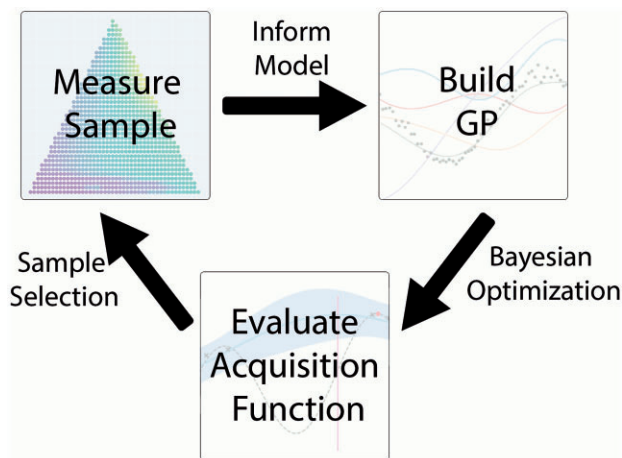


Figure 1: Representation of the active learning (AL) pipeline using GP as the ML surrogate function and Bayesian optimization for materials optimization. The AL framework guides subsequent experiments to hone-in on the target materials.

experiments (see Fig. 1). The active learning scheme identifies a material composition to investigate; pertinent data are then collected from the model to simulate experiment execution; ML is then used to analyze the collected data, and the process repeats.

The past few years have seen a variety of benchmarking studies with a focus on ML model selection [14–19]. More recently, a benchmarking study investigated the choice of active learning scheme for challenges including optimizing 3D-printing parameters for print feature profile, 3D-printed structure parameters for toughness, composite blends for electrical conductivity, molecular mixtures for photovoltaic stability and synthesis parameters for nanoparticle optical absorbance [20]. While not focused on solid-state materials synthesis optimization, this benchmarking study collects datasets that may serve as potential surrogates. Another study presented an exhaustive dataset of solid-state composition and catalytic response [19]. This study investigated the impact of ML model selection and compared performance of one active learning scheme to random selection.

In this work, we present a reference materials dataset of composition, structure and multiple functional properties. The dataset (Fig. 2) is from a fully-characterized thin film library of the Fe-Co-Ni [21] system where structural properties and magnetic properties were mapped across the entire ternary. Datasets from such composition spreads are essential for benchmarking as they contain data for exhaustive experiments across a synthesis parameter space. We use this database to benchmark common off-the-shelf science agnostic BO schemes along with one scientific ML [22] BO scheme for the task of materials optimization. Here, scientific ML refers to ML algorithms with built-in prior physical knowledge. A previous study benchmarked active learning schemes for composition-phase determination on the Fe-Ga-Pd dataset [23]. These benchmark datasets along with others are available at the Resource for Materials Informatics website [24]. A link to the code can also be found in the references [25].

DISCUSSION

A thin film Fe-Co-Ni composition library dataset [21] is used to investigate materials optimization performance of varying active learning schemes. Expert-based data cleaning was performed to simplify benchmark use and a description of these preprocessing steps is given in [Supplemental Information](#). Sample materials span the full ternary composition space of Fe-Co-Ni, with 921 compositions/samples in total. Structure data are collected for

each sample, providing information of the ternary composition–structural phase diagram (Fig. 2a), i.e. how lattice structure varies across the composition space. Two materials properties were measured for each sample—Kerr rotation (Fig. 2b) and magnetic coercivity (Fig. 2c). The Kerr rotation is a measure of magnetization of the material, while the coercivity is a measure of the magnetic hardness [26]. Here the composition dependence of the Kerr rotation is shown to be the less complex of the two, smoothly varying with a broad peak near the binary of Fe_{0.4}Ni_{0.6}. Coercivity shows a broad peak at a similar composition which is overwhelmed by a higher, complex and more ‘granular’ coercivity response near the Co-Ni binary with many local optima. The roughness or complexity of the Kerr rotation and magnetic coercivity were quantified by fitting a two-dimensional GP to both datasets after normalizing (values mapped to 0 to 1) and then investigating the length scales and ratio of variance to length scales. Kerr rotation has length scales of approximately 50% larger than those of magnetic coercivity while both share a similar ratio of variance to length scales. These numbers agree with the visual, qualitative analysis that the Kerr rotation varies far more gradually than the magnetic coercivity.

The list of off-the-shelf active learning acquisition functions [13] benchmarked is presented in Table 1. The off-the-shelf acquisition functions are paired with a GP for the prediction model. Here, μ , σ^2 and Σ are the GP mean, variance and covariance, respectively. $f(x^+)$ is the maximum found in the previous iterations; n is the current iteration number; D is the number of materials to search over; λ is a predefined constant, here set to 0.1; the constants $\sqrt{\ln(Dn^2\pi^2)/3\lambda}$ and $\sqrt{\ln(2n+1)}/8$ balance exploration typified by maximizing σ and exploitation typified by maximizing μ ; N is the probability density function for the normal distribution, Φ and ϕ are the cumulative distribution function and the probability distribution function for the standard normal distribution $N(\mu=0, \sigma=0)$, respectively, ζ is a user-controlled variable for reducing the impact of exploitation and $Unif()$ is the uniform distribution. The off-the-shelf-acquisition functions are paired with a GP using an isometric radial basis function and Gaussian likelihood. These active learning schemes are compared to a recent physics-informed active learning scheme named Closed-loop Autonomous Materials Exploration and Optimization (CAMEO) [4]. CAMEO exploits a fundamental rule in materials science, that a material system’s structure, or phase, is predictive of its functional properties. As such, CAMEO first seeks to maximize knowledge of the target

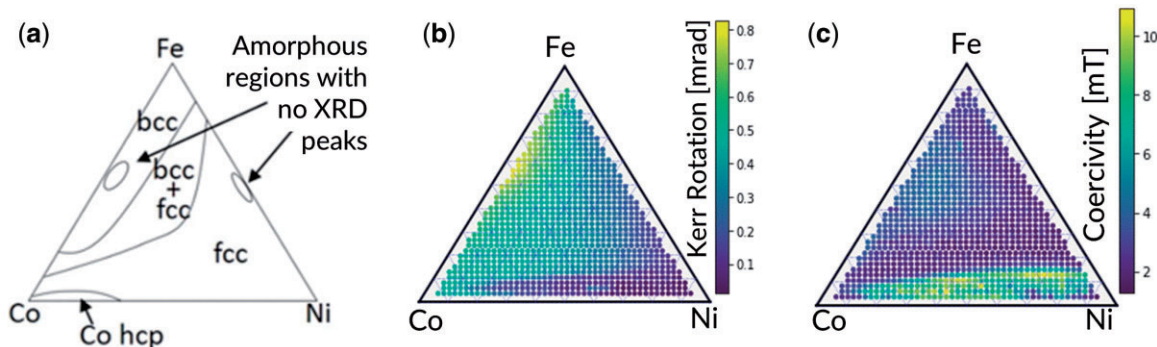


Figure 2: The Fe-Co-Ni library dataset [21]. (a) The composition–structural phase map for the library. This map was determined through expert analysis of collected X-ray diffraction data. The two materials optimization challenges have been chosen for their varying complexity. (b) The Kerr rotation (measured in milliradians) is the least complex dataset. As seen visually, the function is dominated by a large gradually increasing peak with a maximum located along the Fe-Co binary axis. (c) The magnetic coercivity (measured in milli-Tesla) represents a high complexity system with the global maximum and many local maxima located in a small region near the Ni-Co binary axis and one broad small peak located along the Fe-Co binary axis.

Table 1: Acquisition functions

Name	Acquisition function	Description
Exploration	$\operatorname{argmax}_x[\sigma]$	Point where model has max uncertainty
Exploitation	$\operatorname{argmax}_x[\mu]$	Point with predicted max value
Upper Confidence Bound (UCB)	$\operatorname{argmax}_x\left[\mu + \sigma\sqrt{\ln(Dn^2\pi^2)/3\lambda}\right]$	Balance of exploration and exploitation with iteration n dependent balance ratio. As the number of datapoints in a region and the iteration number increase, σ associated with model uncertainty will decrease. The weight on σ increases to maintain exploration.
Add-GP-UCB [27]	$\operatorname{argmax}_x\left[\mu + \sigma\sqrt{\ln(2n+1)/8}\right]$	Balance of exploration and exploitation with iteration-dependent balance ratio. The impact of this ratio is similar to that of UCB (see above).
Thompson sampling	$\operatorname{argmax}_x[f \sim N(\mu, \Sigma)]$	Sample function f from Gaussian distribution with model's predicted mean and covariance. Then identify the point with max value.
Expected improvement	$\operatorname{argmax}_x[(\mu - f(x^+) - \xi)\Phi(Z) + \sigma\phi(Z)]$ $Z = (\mu - f(x^+) - \xi)/\sigma$	Identify point expected to have maximal improvement over past identified maximum
Random sampling	$x^* \sim \operatorname{Unif}(X)$	Sample point at random

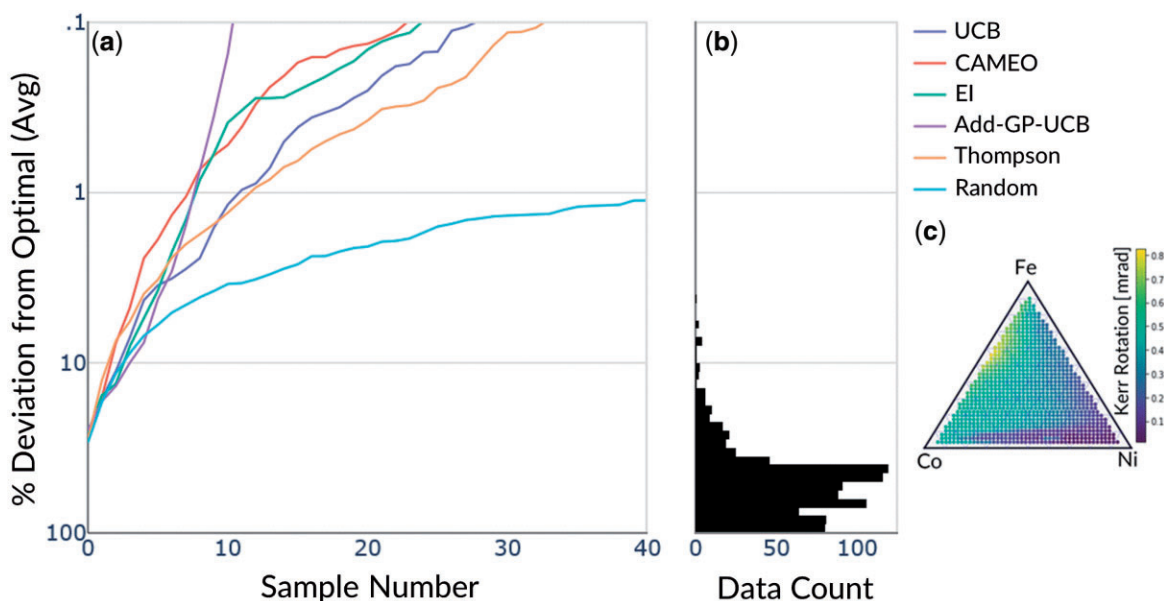


Figure 3: (a) Benchmarking performance of acquisition functions on Fe-Co-Ni Kerr Rotation with reference to random selection. (b) A histogram displaying the counts of data at different percent deviations from the optimal. (c) The Fe-Co-Ni Kerr rotation dataset replotted.

material system's phase diagram. It then uses phase boundaries to segment the search space and guide the search for materials optimization.

The materials optimization process is initialized by selecting one material from the pool of potential Fe-Co-Ni materials with uniform probability, i.e. all materials have equal likelihood of being selected as the starting material. The GP is then used to predict the material property of interest for all materials including those without functional property data. The acquisition function of interest is used to select the next material to investigate. This loop is continued exhaustively for the remainder of the unevaluated data points (materials cannot be selected twice in this implementation) with the same acquisition function. For each acquisition function, the process is repeated 100 times. Performance is computed using minimum regret, defined as:

$$\text{minimum regret} = \max(\text{sampld}) - \max(\text{global}) \quad (1)$$

The mean and mean confidence intervals for the minimum regret are computed over the 100 runs, and the mean is plotted in Figs 3 and 4 for the Kerr rotation and coercivity, respectively. Figures with both mean and the confidence intervals plotted are available in Supplemental Information.

For the lower complexity Fe-Co-Ni Kerr rotation dataset, locating the one dominant peak can be achieved with a simple estimated gradient ascent method [7]. As a result, greedier acquisition functions that focus on exploitation have better performance. Nevertheless, all acquisition functions reach within 0.1% deviation from the maximum within 5% of the 921 data points. Add-GP-UCB [27] reaches the goal within approximately 11 samples; about double the speed of its next competitor. In second place is the physics-informed CAMEO algorithm which must expend initial iterations to identify the phase diagram. This forced exploration puts CAMEO at a disadvantage when more aggressive exploitation provides better performance.




Figure 4: (a) Benchmarking performance of acquisition functions on Fe-Co-Ni magnetic coercivity dataset with reference to random selection. (b) Adjacent to the performance graph is a histogram displaying the counts of data at different percent deviations from the optimal. (c) The dataset is complex, with a broad, minor peak near the Fe-Co binary and high roughness with many local maxima near the global maximum close to the Co-Ni binary.

All acquisition functions perform significantly better than random sampling due to the simplicity of the target function. If the data are known beforehand to have gradual changes in intensity and one dominant peak, an exploitation-focused acquisition algorithm is preferable.

The high-complexity material dataset of Fe-Co-Ni magnetic coercivity has a large broad maximum near the Fe-Co binary and high roughness/variation with many local maxima in the proximity of the global maximum, all of which can serve to distract an acquisition function. The high roughness near the global maximum results in a small indicative composition region, making the maximum difficult to find. The majority of acquisition functions perform poorly, easily becoming stuck in local maxima. Expected Improvement performs well in comparison to other common acquisition functions, and is more exploratory than the other methods, while still outperforming random. While the physics-informed CAMEO algorithm is capable of narrowing the search space to those phase regions that promise to hold a maximum, the high roughness of the target phase regions still manages to distract CAMEO from the global maximum. This is likely due to the use of the UCB algorithm once CAMEO switches from phase mapping to materials optimization. The number of samples required to reach 0.1% from the optimum is larger than that of the simpler Kerr rotation challenge.

The Fe-Co-Ni Kerr dataset of composition, X-ray diffraction, Kerr rotation and magnetic coercivity provides some interesting insights into active learning use for real-world tasks materials optimization. When a smooth, simple landscape is identified, schemes with greater focus on exploitation succeed. The alternative is also true with exploration being preferable with highly complex landscapes. This suggests an active learning scheme that is iteration dependent. At each iteration, one can compare prediction accuracy and quantify landscape complexity to determine the preferred acquisition function for the next iteration, shifting the balance between exploration and exploitation. Nevertheless, expected improvement is the overall winner,

performing well for either task, despite the complexity of the coercivity data.

The physics-informed CAMEO method lags in performance when the challenge is simple, due to its forced steps of iteration—it must first converge on a phase diagram before seeking an optimal material. It is also able to narrow the search space with a highly complex landscape, but the current choice of UCB for optimization within the target phase region performs poorly and should potentially be replaced by an alternative method such as expected improvement. The authors hope that the Fe-Co-Ni and Fe-Ga-Pd [4] datasets can spur interest in benchmarking and developing novel active learning schemes for real-world challenges of materials exploration and optimization.

SUPPLEMENTARY DATA

Supplementary data is available at *Oxford Open Materials Science Journal* online.

FUNDING

This work was funded by NIST Cooperative Agreement 70NANB17H301 and an Office of Naval Research MURI through grant #N00014-17-1-2661.

AUTHORS' CONTRIBUTIONS

A.G.K., I.T., H.L. and A.M. contributed to conceptualization; A.G.K., A.W. and A.M. contributed to methodology; A.W. and A.G.K. contributed to software; A.G.K. contributed to validation; A.W. contributed to formal analysis; A.W. contributed to investigation; A.G.K. and I.T. contributed to resources; A.W. and A.G.K. contributed to data curation; A.W. contributed to writing—original draft preparation; A.G.K. and I.T. contributed to writing—review and editing; A.W. and A.G.K. contributed to visualization; A.G.K. and I.T. contributed to supervision; A.G.K. and I.T. contributed to funding acquisition.

CONFLICT OF INTEREST STATEMENT

The authors have no conflicts of interest to declare.

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.

REFERENCES

- Schilling A, Cantoni M, Guo JD *et al.* Superconductivity above 130 k in the hg–ba–ca–cu–o system. *Nature* 1993;363:56–8.
- Green ML, Takeuchi I, Hattrick-Simpers JR. Applications of high throughput (combinatorial) methodologies to electronic, magnetic, optical, and energy-related materials. *J Appl Phys* 2013;113:231101.
- Potyralo R, Rajan K, Stoewe K *et al.* Combinatorial and high-throughput screening of materials libraries: review of state of the art. *ACS combinatorial science*. 13 (6) 2011. 579–633. <https://doi.org/10.1021/co200007w>.
- Kusne AG, Yu H, Wu C *et al.* On-the-fly closed-loop materials discovery via Bayesian active learning. *Nat Commun* 2020;11:1–11.
- Kusne AG, Gao T, Mehta A *et al.* On-the-fly machine-learning for high-throughput experiments: search for rare-earth-free permanent magnets. *Sci Rep* 2014;4:6367.
- Xue D, Balachandran PV, Hogden J *et al.* Accelerated search for materials with targeted properties by adaptive design. *Nat Commun* 2016;7:11241.
- NIST/SEMATECH e-Handbook of Statistical Methods. <https://www.itl.nist.gov/div898/handbook/>, 2022.
- Braham EJ, Davidson RD, Al-Hashimi M *et al.* Navigating the design space of inorganic materials synthesis using statistical methods and machine learning. *Dalton Trans* 2020;49:11480–8.
- Stach E, DeCost B, Kusne AG *et al.* Autonomous experimentation systems for materials development: a community perspective. *Matter* 2021;4:2702–26.
- Epps RW, Bowen MS, Volk AA *et al.* Artificial chemist: an autonomous quantum dot synthesis bot. *Adv Mater* 2020;32:2001626.
- MacLeod BP, Parlane FGL, Morrissey TD *et al.* Self-driving laboratory for accelerated discovery of thin-film materials. *Sci Adv* 2020;6:eaaz8867.
- Rasmussen CE Gaussian processes in machine learning. In: Bousquet O, von Luxburg U, Rätsch G. (eds) *Advanced Lectures on Machine Learning*. ML 2003. Lecture Notes in Computer Science, Vol. 3176, Springer, Berlin, Heidelberg. https://doi.org/10.1007/978-3-540-28650-9_4
- Shahriari B, Swersky K, Wang Z *et al.* Taking the human out of the loop: a review of Bayesian optimization. *Proc IEEE* 2016; 104:148–75.
- Fung V, Zhang J, Juarez E *et al.* Benchmarking graph neural networks for materials chemistry. *NPJ Comput Mater* 2021;7:1–8.
- Dunn A, Wang Q, Ganose A *et al.* Benchmarking materials property prediction methods: the Matbench test set and Automatminer reference algorithm. *NPJ Comput Mater* 2020;6:1–10.
- Clement CL, Kauwe SK, Sparks TD. Benchmark AFLOW data sets for machine learning. *Integr Mater Manuf Innov* 2020;9:153–6.
- Breuck P-PD, Evans ML, Rignanese G-M. Robust model benchmarking and bias-imbalance in data-driven materials science: a case study on MODNet. *J Phys Condens Matter* 2021; 33:404002.
- Thiyagalangam J, Shankar M, Fox G *et al.* Scientific machine learning benchmarks. *Nat Rev Phys* 2022;4:413–20.
- Rohr B, Stein HS, Guevarra D *et al.* Benchmarking the acceleration of materials discovery by sequential learning. *Chem Sci* 2020;11:2696–706.
- Liang Q, Gongora AE, Ren Z *et al.* Benchmarking the performance of Bayesian optimization across multiple experimental materials science domains. *NPJ Comput Mater* 2021;7:1–10.
- Yoo YK, Xue Q, Chu YS *et al.* Identification of amorphous phases in the Fe–Ni–Co ternary alloy system using continuous phase diagram material chips. *Intermetallics* 2006;14:241–7.
- Baker N, Frank A, Timo B *et al.* *Workshop Report on Basic Research Needs for Scientific Machine Learning: Core Technologies for Artificial Intelligence*, 2019. <https://www.osti.gov/biblio/1478744/>. doi:10.2172/1478744.
- Kusne AG, McDannald A, DeCost B *et al.* Physics in the machine: integrating physical knowledge in autonomous phase-mapping. *Front Phys* 2021;10:815863.
- Remi: Resource for Materials Informatics. Home - REMI. <https://pages.nist.gov/remi/>, Version 0.1.
- Wang A. *Starter code for benchmarking active learning methods*. GitHub. <https://github.com/alex-aa-wang/Benchmarking-Active-Learning-Starter>, 2022
- O’Handley RC. *Modern Magnetic Materials: Principles and Applications*. New York: Wiley, 2000.
- Kandasamy K, Schneider J, Póczos B. High dimensional Bayesian optimisation and bandits via additive models. In: *International Conference on Machine Learning*, pp. 295–304, Lille: PMLR, 2015.