

Autonomous Materials Exploration Integrates Automated Phase Identification and AI Agents Enhanced by Human Guidance

Ming-Chiang Chang^{1,*}, Maximilian Amsler^{1,*}, Duncan R. Sutherland¹, Sebastian Ament², Katie R. Gann¹,
Lan Zhou,^{3,4} Louisa M. Smieska⁵, Arthur R. Woll⁵, John M. Gregoire^{3,4}, Carla P. Gomes²,
R. Bruce van Dover¹ and Michael O. Thompson^{1,§}

¹Department of Materials Science and Engineering, *Cornell University*, Ithaca, New York 14853, USA

²Department of Computer Science, *Cornell University*, Ithaca, New York 14853, USA

³Division of Engineering and Applied Science, *California Institute of Technology*, Pasadena, California 91125, USA

⁴Liquid Sunlight Alliance, *California Institute of Technology*, Pasadena, California 91125, USA

⁵*Cornell High Energy Synchrotron Source*, *Cornell University*, Ithaca, New York 14853, USA



(Received 17 April 2026; accepted 22 July 2026; published 10 September 2026)

Autonomous experimentation holds the potential to accelerate materials development by combining artificial intelligence (AI) with modular robotic platforms to explore extensive combinatorial chemical and processing spaces. Such self-driving laboratories can not only increase the throughput of repetitive experiments, but also incorporate human domain expertise to drive the search toward user-defined objectives, including improved materials performance metrics. We present an autonomous materials synthesis extension to SARA, the Scientific Autonomous Reasoning Agent, utilizing phase information provided by an automated probabilistic phase labeling algorithm to expedite the search for targeted phase regions. By incorporating human input into an expanded SARA-H (SARA with human-in-the-loop) framework, we enhance the efficiency of the underlying reasoning process. Using synthetic benchmarks, we demonstrate the efficiency of our AI implementation and show that the human input can contribute to significant improvement in sampling efficiency. We conduct experimental active learning campaigns using robotic processing of thin-film samples of several oxide material systems, including Bi_2O_3 , SnO_x , and Bi-Ti-O , using lateral-gradient laser spike annealing to synthesize and kinetically trap metastable phases. We showcase the utility of human-in-the-loop autonomous experimentation for the Bi-Ti-O system, where we identify extensive processing domains that stabilize $\delta\text{-Bi}_2\text{O}_3$ and $\text{Bi}_2\text{Ti}_2\text{O}_7$, explore dwell-dependent ternary oxide phase behavior, and provide evidence confirming predictions that cationic substitutional doping of TiO_2 with Bi inhibits the unfavorable transformation of the metastable anatase to the ground-state rutile phase. The autonomous methods we have developed enable the discovery of new materials and new understanding of materials synthesis and properties.

DOI: [10.1103/c362-349b](https://doi.org/10.1103/c362-349b)

I. INTRODUCTION

Technological developments to address global challenges are critically dependent on the deployment of new materials with enhanced properties, particularly today in the realm of clean power generation and energy storage [1,2]. Although the Edisonian method of trial and error has succeeded in the discovery of many materials in the past, the advent of innovative synthesis techniques, accompanied by the added dimension of chemical complexity in the search for new

materials, has opened up vast processing spaces that require detailed exploration. High-throughput (HT) synthesis [3,4] has emerged as a powerful approach to combinatorially map large portions of this search space, supported by accurate *ab initio* calculations [5,6] that predict novel materials using automated computational screening [7,8] along with advanced global optimization algorithms [9–11]. The large data sets that can now be produced with such HT methods require equally advanced data analysis techniques to elucidate complex processing-composition-structure-property relationships. Simultaneously, curated materials information has been accumulated in numerous online repositories accessible for further analysis [5,6,12], marking a significant advancement in data democratization efforts that may also accelerate materials development [13].

Data from these direct and *in silico* experiments are also critical as training sets for many machine learning (ML) models, improving predictive models to guide material development. Prominent recent examples include AlphaFold [14], which relies on a library of protein data to build models of protein folding, and GNoME [15], an inorganic

*These authors contributed equally to this work.

†Contact author: mc2663@cornell.edu

‡Contact author: amsler.max@gmail.com

§Contact author: mot1@cornell.edu

Published by the American Physical Society under the terms of the [Creative Commons Attribution 4.0 International](https://creativecommons.org/licenses/by/4.0/) license. Further distribution of this work must maintain attribution to the author(s) and the published article's title, journal citation, and DOI.

materials generation framework based on graph neural networks trained on materials data. While pretrained ML models like deep neural networks and foundation models [16–18] excel at interpolation and can work well in the presence of sufficiently large data, active learning (AL) approaches are better suited to guide an exploratory process to gain insight when existing data are sparse and hence are more useful for accelerating material experiments [19]. AL algorithms aim to successively improve models (ML or others) by proposing optimal experiments that provide the most useful information given existing knowledge, thereby driving the learning process effectively with a feedback loop and ensuring data and time efficiency. Successful demonstrations of AL in materials discovery have been reported for several applications, including phase-change materials [20], NiTi-based shape memory alloys with low thermal hysteresis [21], and selective synthesis of carbon nanotubes [22,23]. Meanwhile, several autonomous robotic laboratories for closed-loop synthesis of novel organic compounds have been established [24–28], while analogous infrastructures for inorganic materials synthesis are still under development [29–31].

Despite the enormous potential that AI offers for autonomous and efficient experimental planning in high-dimensional space [32–34], physical reasoning with experimental results remains a challenge for AI. In contrast, humans struggle to understand high-dimensional distributions but are able to interpret results using intricate physical intuition and past experiences to guide experiments. For example, when humans observe unexpected behavior, they often significantly alter what had been identified as the next experiment. Ideally, experimental workflows should leverage the unique advantages of AI and humans: Autonomous AI agents can readily relieve scientists from the details of experiment planning, allowing them to instead focus on drawing physical meaning from the data in real time. This is akin to AI assistants, such as agentic coding in the programming domain, where the AI assists in writing code for very specific functions while engineers concentrate on the overall code integration and high-level tasks [35–37]. The critical question is where the appropriate boundary should be established, and how to create an effective interface to facilitate a complex, well-orchestrated interplay between humans and machines, ultimately optimizing the “human-in-the-loop” AI for complex material discoveries [38,39].

To this end, we expand the capabilities of SARA, the Scientific Autonomous Reasoning Agent [40], with an enhanced workflow that incorporates human scientist reasoning that directs evolving objectives within an autonomous experimental campaign. SARA employs principled reasoning grounded on physical laws and uncertainty estimates. In contrast to current large language model (LLM) efforts, SARA’s agentic aspects incorporate ML and interactions with tools such as modeling resources, experimental synthesis, and analytical characterization.

In particular, we demonstrate “SARA with human-in-the-loop” (SARA-H) in campaigns that integrate autonomous high-throughput processing with lateral-gradient laser spike annealing (lg-LSA) [41] and real-time structural characterization [42] of thin films. As the knowledge of the processing-composition landscape develops, the human in the

loop alters priorities, moving from exploration to optimizing synthesis conditions for targeted phases (exploitation). The advantage of incorporating phase information and advanced acquisition functions is demonstrated in both synthetic tests and AL experiments conducted on three complex oxide systems, Bi_2O_3 , SnO_x suboxides, and the pseudobinary Bi-Ti-O system. We discuss our findings for Bi_2O_3 and Bi-Ti-O extensively in Secs. IIC 1 and IIC 2. The SnO_x suboxide system serves as a reference benchmark and the details are included in the Supplemental Material [43]. By incorporating human expert interaction within the AL cycles, we enable the targeted synthesis of specific material phases. For example, within the Bi-Ti-O system, we targeted and mapped the phase field of bismuth-substituted anatase, a phase that has been predicted to exhibit enhanced photocatalytic behavior [51–53]. With these advancements, SARA-H represents a significant milestone in the development of accelerated experimental materials discovery methods.

II. RESULTS AND DISCUSSION

A. Targeted material synthesis with SARA-H

The SARA-H framework was designed to efficiently explore synthesis phase diagrams and to target the synthesis of specific phases. We particularly target metal oxides that can form numerous metastable polymorphs under ultrafast annealing conditions achieved with lg-LSA. Typically, the conditions under which phases can be synthesized are constrained by the precursors, stoichiometry, synthesis methods, and/or limited ranges of thermodynamic processing parameters including temperature and pressure. Even if synthesized and quenched to ambient conditions, many phases often decompose into more stable (but less desirable) states upon reheating or under operating temperatures in relevant device applications. Therefore, significant efforts have been made to better understand the thermodynamic scale of metastability in inorganic materials [54,55] and to develop strategies to stabilize them under operating conditions after formation [56].

The general workflow of SARA-H is shown in Fig. 1. Our experimental synthesis module processes amorphous thin-film compositional libraries, which were previously prepared by depositing precursor thin films using radio-frequency (rf) magnetron sputtering onto heavily doped silicon wafers with a thermally grown silicon oxide buffer layer. Composition spreads can be deposited by cosputtering multiple targets to achieve a smoothly varying compositional gradient across the wafer, with the stoichiometry being a function of position on the wafer, $\mathbf{c}(\mathbf{x}, \mathbf{y})$.

The processing of these thin-film libraries was conducted using lg-LSA, which forms both equilibrium and metastable phases at high temperatures, and kinetically traps them to ambient conditions by rapid heating and quenching. In contrast to other traditional annealing methods for thin film samples, including hot plates, furnaces, or rapid thermal annealing [57], lg-LSA enables controlled rapid processing in spatially confined regions, with quench rates of 10^4 – 10^7 K/s and peak temperatures T_p limited only by the melting point of the substrate (1410 °C for Si substrates). The scanned laser has an intentional lateral intensity profile (typically near




FIG. 1. The overall workflow of SARA-H. (a) The autonomous experimental workflow of SARA-H. The leftmost panel represents SARA-H's synthesis technique, which is based on the annealing of amorphous thin-film samples with lg-LSA under specific conditions. A 3D-rendered experimental setup is shown, with the laser to the left, a camera to the right representing characterization and the power profile in the backdrop, and the thin-film sample mounted on a stage. An annealed lg-LSA stripe (second panel) is then characterized with narrow-beam synchrotron x-ray diffraction. The top image shows a micrograph of an annealed stripe, together with the nominal XRD footprint that is scanned across the stripe, while the bottom image presents the corresponding XRD heat map. The third panel illustrates analysis of the diffraction data on the fly using the probabilistic multiphase-labeling algorithm CrystalShift. Finally, the expected phase fraction from the probabilistic phase model can be utilized to construct a phase diagram (shown as a probability distribution in a high-dimensional processing space), which is progressively improved using Bayesian optimization. The AL loop is closed by performing the next optimal synthesis. (b) The optimization strategy may be augmented with human expertise to further enhance the targeted synthesis of specific phases. For a given material system, the expert prepares a list of candidate phases to study in the active learning run. During the run, the expert monitors results and determines which candidate phases to retain in the active learning loop for optimization. This human-in-the-loop setup creates an interface for experts to control the behavior of the autonomous experimentation. Details on the time scales of every step involved in the SARA-H workflow are compiled in Sec. F of the Supplemental Material [43].

Gaussian), giving each lg-LSA stripe a spatially varying thermal profile. Process temperature and time are dependent on the laser power and the scan velocity, which are characterized by a peak temperature at the center of the stripe T_c and a dwell time related to the heating and quench rates. For the experiments in this work, the full width at half maximum (FWHM) of the temperature distribution was $x_w \approx 360 \mu\text{m}$ across the stripe, with some variation based on the anneal conditions. The inset in the *experimental synthesis* panel of Fig. 1(a) is a schematic illustrating the annealing temperature distribution during lg-LSA. The heating durations and the associated heating/quench rates are characterized by a dwell

time τ defined as the FWHM of the laser spot divided by the stage velocity. The cooling/quench rate, in particular, can be extremely high (up to 10^7 K/s) due to direct contact of the thin film with a thermally conductive Si substrate. Each lg-LSA stripe has a common dwell τ , but a continuous range of temperature conditions across the stripe, yielding multiple phases as a function of the peak annealing temperature.

Following each experimental lg-LSA stripe anneal, the resulting phase behavior was analyzed using multiple appropriate characterization techniques, as illustrated by the *anneal characterization* panel in Fig. 1(a). In previous work [58], we employed microscopy imaging and optical spectrometry to

identify structural phase changes that were directly correlated with alterations in the optical properties of oxide films [59]. These optical measurements across each Ig-LSA stripe were subsequently used to delineate phase boundaries without explicit crystallographic phase identification, resulting in the creation of an unlabeled processing phase diagram. In the current work, we modified analysis agents in SARA-H to incorporate powder x-ray diffraction (XRD) to directly inform the agents of specific crystalline phases synthesized during the process, as shown in the *probabilistic phase analysis* panel in Fig. 1(a). To achieve this, a focused 20 μm wide x-ray beam was scanned across each stripe, generating diffraction patterns as a function of lateral location (associated with the temperature gradient). The complete diffraction data of an annealed stripe were fed into a non-negative matrix factorization (NMF) algorithm [60] to reduce the dimensionality and accelerate data processing. This specific NMF algorithm identified the most representative XRD patterns that could be used to reconstruct the full XRD data matrix with minimal error, thereby ensuring that it did not hallucinate artificial XRD patterns. To label the phases, we utilized CrystalShift [42], a tree-based probabilistic algorithm designed for rapid multiphase labeling. A high-level overview of the CrystalShift workflow is provided in Sec. IV. The identified XRD bases were fed into this algorithm to generate the final labels. The labeling results were then combined with the linear coefficients of the NMF to provide an expected phase activation distribution, which serves as a robust proxy for the phase fraction and is defined in Eq. (6), along with their related uncertainties as a function of temperature (for the given dwell). This is illustrated by the *phase diagram* panel in Fig. 1(a).

The expected phase activation of the annealed stripes was used to create and continuously update a Gaussian process (GP) model of a phase map within the parameters $\{\mathbf{c}, \tau, T_p\}$, which encodes the fraction p_k of each present phase k . This model not only aids in constructing an accurate phase diagram, but also doubles as the surrogate model for optimizing the acquisition functions for AL cycles, which determines the next anneal based on the available information. Although the previous version of SARA was limited to optimizing the phase-boundary map by targeting maximal gradient regions in optical data [58], the probabilistic phase distribution produced by XRD and CrystalShift allows more refined objectives. In particular, we can employ Bayesian optimization to perform targeted materials synthesis. At the core of such optimizations, we utilized a recently developed acquisition function based on expected improvement (EI) called logEI [61], which prevents numerical underflow and is therefore more stable and easier to optimize than the regular EI function; this function then served as the foundation for our efficient sampling strategies. During both surrogate construction and the optimization of the acquisition function, proper propagation of experimental errors was included to increase the robustness and performance of the AL agent. A comprehensive benchmark of different sampling methods is discussed in Sec. IIB.

Finally, we incorporate human-machine interaction with SARA-H using a simple yet powerful interface, represented by the *expert interaction* panel in Fig. 1(a). The advantage of human scientists involved in the autonomous experiment

is that they can provide the model with higher-level, non-deterministic information that enhances the decision-making process based on reviewing the current state of the model and the aggregated data. The types of expert knowledge can be diverse, ranging from visually identifying the melting point or delamination conditions of a specific phase, to constraining the search boundaries of the model or correcting misclassifications of XRD data.

In this work, human experts can modify the objective of the AL cycle on the fly. At the start of an AL trial, the analysis agent is typically provided with a complete list of potential phases in the target materials system, shown as the *cycle sampling* tab in Fig. 1(b). The AL cycles through this list and attempts to maximize the expected phase fraction of one targeted phase in each loop. The human can examine the data during the run and decide to add or remove candidate phase(s) [illustrated by the *cycle over subset* tab in Fig. 1(b)]. The reasons for removing candidate phases may include, but are not limited to, the sufficient convergence of the phase field for a specific phase, the likelihood that the phase cannot be stabilized, or a greater interest in certain phases over others. This round-robin fashion cycle sampling method prevents the expensive evaluation of a multivariate Gaussian process required for optimizing entropylike objectives.

This implementation of SARA-H is also capable of mimicking a scientist's behavior when searching for a specific phase and attempting to optimize the conversion rate of the same phase for every loop [denoted as *target sampling* in Fig. 1(b)]. Our interface enables a spectrum of acquisition schemes across the exploration-exploitation extrema by adjusting candidate phases, allowing real-time control over autonomous experimental behavior. In particular, we incorporate this human-AI collaboration scheme in the autonomous experimental process for the complex pseudobinary Bi-Ti-O system (discussed in detail in Sec. IIC2).

B. Synthetic benchmarks for sampling strategies

In AL and Bayesian optimization strategies, we use a sampling framework wherein the underlying selection criterion shifts from a uniform random selection (used in random sampling) to dynamically updated acquisition functions, for example, an uncertainty sampling acquisition function or a logEI acquisition function. Throughout the remainder of this article, any reference to these strategies implicitly denotes this generalized sampling perspective. For comparisons of the acquisition function and objective between the different strategies, see Table S1 in the Supplemental Material [43].

The task of identifying the optimal synthesis conditions for a specific phase differs from the common problems typically used to benchmark AL methods, such as the Ackley function [62], which are often designed to have challenging aspects for AL (e.g., multiple local minima, narrow valleys, flat loss contours, etc.). In efforts to identify the experimental condition that contributes to phase formation rates, the target distribution is generally assumed to be straightforward, featuring a well-defined maximum and nonzero values contained within a single continuous volume in the parameter space, with few exceptions where phases may exist in more than one volume. To quantify the benefits of using phase information




FIG. 2. Performance of various active learning sampling strategies in a synthetic phase space. (a) The median regret on one of the phases for a materials system containing a total of eight phases against the number of iterations with different sampling strategies. Solid lines are the stripe variant while the dashed lines represent their single-point counterpart. (b) The IQR of the regret versus the number of iterations for each sampling strategy, for the same phase as in panel (a). Compared to random sampling, the regret IQR for other methods rises and decreases earlier and more sharply. This indicates the efficient exploration and exploitation of active learning methods. (c) The enhancement factors represent the ratio between the objective values obtained from AL and those achieved through random sampling at various stages, averaged over all eight phases present in the materials system. This quantity illustrates the degree to which the results improve when utilizing AL techniques. (d) The acceleration factors represent the ratios of the average iterations needed to achieve the same objective value for random sampling with respect to each AL strategy, averaged over all eight phases that exist.

in identifying optimal synthesis conditions, and to statistically compare different sampling strategies, we conducted tests on an experimental dataset in a three-dimensional (3D) parameter space, where the first two degrees of freedom (DOF) represent the peak anneal temperature and dwell time, and the third DOF represents a compositional variable. The system contains eight phase regions that were fitted from real experimental data obtained for Bi-Ti-O, a system that will be discussed in more detail in Sec. IIC 2.

We assessed the efficiency of various sampling strategies primarily using two benchmark metrics: the median and the interquartile range (IQR) of regret r , where r is a measure that expresses the relative difference between the maximum phase activation of phase k found by the AL model, $a_k(\theta^*)$, and the ground truth maximum phase activation, $a_k(\theta)$, under a given set of parameters $\theta = (\mathbf{c}, \tau, T)$. The median regret is its median for all test trials at a specific iteration.

We benchmarked four different sampling strategies, including random sampling and uncertainty sampling as

baseline methods, and the two methods proposed in this work, namely, target sampling and cycle sampling (results for pseudorandom methods are shown in Fig. S1 of the Supplemental Material [43]).

For each of the above AL sampling strategies, their stripe acquisition variants were also tested, a method developed and introduced in our previous work [58]. The stripe acquisition considers a temperature *range* of an lg-LSA stripe to identify the optimal conditions for the subsequent anneal, which can be considered a form of batch processing. Benchmarking the stripe variants evaluates the extent to which the temperature-gradient phase information in lg-LSA stripes enhances the performance of our workflow. 500 synthetic AL trials were conducted for each acquisition scheme to acquire performance statistics. The stripe variants were repeated 100 times due to their longer run times.

Figure 2(a) shows the representative median regret result for one of the eight targeted phases. Our results clearly demonstrate that the use of AL techniques significantly

reduces regret compared to random sampling, although not all of these techniques perform equally well.

The uncertainty sampling strategy served as a baseline for using the active learning scheme. It selects the condition for the next AL iteration that maximally reduces the overall uncertainty of the model without knowing any phase information. It outperforms random sampling in the long run but its ability to identify the best condition for the target phase plateaus due to its inability to *exploit* accumulated knowledge of the system.

Target sampling employs the logEI acquisition function to maximize the expected phase fraction of a single, targeted phase. This allows the AL to hone in on one specific phase and is, by design, better suited to exploit phase information compared to uncertainty sampling, which samples the entire parameter space more uniformly.

In contrast, cycle sampling iterates through all possible candidate phases (12 in total, with 8 ground-truth phases and 4 distractor phases), with the aim of optimizing the activation of a single phase at each AL iteration. This strategy mimics a scientist who has a list of potentially interesting phases but is not yet certain which phase or subset of phases to explore in greater detail. Although we found this strategy to be less efficient for identifying an optimal synthesis condition for one specific phase compared to uncertainty sampling in the early stages of an AL campaign, it generally outperformed uncertainty sampling after several AL iterations, showing the gain from incorporating phase information into decision making. The utilization and exploitation of phase-aggregated information is particularly evident in cycle sampling, as clear steps reducing the regret can be observed in the dashed green line. The iterations at which these steps occur correspond to the iterations at which the phase plotted in Fig. 2(a) is the target phase.

All sampling strategies benefit significantly when the stripe acquisition method was incorporated, illustrated by the difference between the dashed and solid lines of the same color in Fig. 2(a). This demonstrates effective integration of the high-throughput characterization of lg-LSA into the acquisition function. It is an example of the general value of incorporating dense and relevant information from high-throughput experiments into AL.

Next, we investigated the IQR of regret to better understand the trade-off between exploration and exploitation in the various AL strategies, as shown in Fig. 2(b) for the same targeted phase as in Fig. 2(a). The IQR is defined as the difference between the 75th and 25th percentiles of the regret distribution in the statistical ensemble, and can be used as a proxy of the behavior of AL strategies. Large IQR values arise when each statistical AL run samples different regions of the parameter space (exploration), while small IQR values indicate similar behaviors among each AL instance (exploitation). In Fig. 2(b), all IQR curves for every sampling strategy follow a similar trend: Initially, the IQR increases, then experiences a downturn, and eventually flattens out gradually. These three domains can be interpreted as follows.

First, the initial increase in the IQR indicates that the AL is exploring the space, which leads to significantly different maximum phase activation observed for each AL trial. The steeper this increase, the more efficient the exploration phase.

Stripe sampling variants are particularly efficient since more information is collected from each stripe.

Second, after the exploration stage, all methods experience a downturn in IQR within the second domain. The sources for this decrease are twofold: (1) it may be caused by the exploitation of information by the AL runs. The drop in IQR indicates that AL instances have explored enough of the parameter space to take advantage of the information to optimize the phase and have begun to reach a consensus on where the target phase should be with a final, ultimately similar $a_k(\theta^*)$, which leads to a smaller IQR. (2) There is an effect of overlapping information in the newly sampled point with the previous one, which is, in fact, the only source that leads to a lower IQR in random sampling. The steeper the drop in IQR, the more efficiently the AL is exploiting the phase-space information. Using the above interpretation, we observe that the IQR curves of cycling and target sampling overlap with that of uncertainty sampling at early iterations, as they are mathematically equivalent before any significant phase fraction has been discovered. However, after the IQR starts to decrease, the target sampling shows a much larger negative slope. This observation, along with lower regret compared to uncertainty sampling, demonstrates that target sampling strategies exhibit markedly faster learning compared to uncertainty or random sampling.

Third, within the last domain, the IQR curves plateau at a residual value. A large residual IQR indicates that additional AL instances do not converge to the same final model. A small residual spread is critical as it provides evidence for both efficient learning and robust model convergence. Note that a similar behavior is observed for the IQR curves of all eight phases in our benchmark set, as shown in Fig. S2 of the Supplemental Material [43].

Overall, our synthetic benchmarks demonstrate that the target and cycle sampling strategies perform best, significantly outpacing both random sampling and uncertainty sampling. Note that if the candidate set is reduced due to human intervention, the sampling performance is expected to be between the two extrema, namely, the target sampling and the cycle sampling.

To quantify the benefit of using AL, the enhancement and acceleration in the maximum activation found were benchmarked and are shown in Figs. 2(c) and 2(d), respectively, using the metric proposed in Liang *et al.* [63]. The results show that AL techniques can accelerate the process of finding the best condition for synthesis of a specific phase by a factor of 2–5 and can find a conversion rate that is roughly two times higher than random sampling. Further discussions on both the enhancement and acceleration factors are provided in Sec. B of the Supplemental Material [43].

C. Active learning experiments in real materials systems

1. Active learning on unary systems

Having established the effectiveness of cycle sampling based on synthetic tests, we turn our attention to the synthesis of real materials to explore phase formation in complex oxides. We first considered the Bi_2O_3 system, which has a phase diagram encompassing at least five distinct experimentally observed polymorphs [64,65]. The metastable $\delta\text{-Bi}_2\text{O}_3$ phase




FIG. 3. (a) Processing phase diagrams of the Bi_2O_3 system after SARA-H's converged AL cycles with 95 iterations. Blue and purple domains represent the presence of $\beta\text{-Bi}_2\text{O}_3$ and $\delta\text{-Bi}_2\text{O}_3$, respectively. The saturation of each region is proportional to the corresponding phase activation with amorphous domain represented as uncolored regions. The dashed black lines approximate the phase boundaries and are intended as a guide for the eye. (b) R_s^2 convergence as a function of the AL iterations for each of the two observed phases, with their average shown as the dashed line.

has attracted considerable attention, especially in the context of its application as a solid electrolyte in fuel cells [66]. The defective fluorite-type crystal structure of the $\delta\text{-Bi}_2\text{O}_3$ phase is characterized by a high concentration of oxygen vacancies, leading to the highest oxygen ion conductivity among solid oxides documented to date [67]. Its application is hindered by a narrow thermodynamic stability window, limited to temperatures between 727 and 824 °C, preventing its use on an industrial scale. Attempts to stabilize $\delta\text{-Bi}_2\text{O}_3$ at room temperature through yttrium [68] or rare earth metal substitutions result in compromised ion conductivity. Consequently, ongoing efforts focus on identifying pathways to maintain phase-pure $\delta\text{-Bi}_2\text{O}_3$ at ambient conditions using grain boundary engineering [69] and sophisticated thermal processing [70].

We previously utilized SARA to autonomously explore the phase map of $\delta\text{-Bi}_2\text{O}_3$ using lg-LSA [58], identifying the boundaries of the sought after $\delta\text{-Bi}_2\text{O}_3$ within the processing phase diagram using purely optical techniques, including microscopy and reflectance spectrometry. In this study, we reinvestigated this system, now constructing a fully labeled phase diagram on the fly using CrystalShift's capability to rapidly disentangle XRD patterns. The Bi_2O_3 system is particularly well suited as a use case for our new algorithm, allowing for the comparison of the resulting phase diagrams.

For SARA-H's closed-loop experiment, we began with a pristine thin film of amorphous Bi_2O_3 , deposited under the same conditions as in our prior work [58]. The AL agents were initialized with simulated XRD patterns derived from a selection of atomic structures of the potential candidate phases, specifically the α , β , γ , δ , and ϵ phases of Bi_2O_3 retrieved from ICSD [12], utilizing cycle sampling for all phases in the candidate pool. As described in detail in Sec. II B, this sampling technique aims to iteratively maximize the phase activation of a phase $k \in \{\alpha, \beta, \gamma, \delta, \epsilon\}$, using the logEI algorithm to determine the next experiment [61]. In SARA-H, the relevant phase k is cycled over the can-

didate phases at each AL iteration; hence, in every step, a different phase is targeted. Note that although only one phase is targeted at each iteration, the knowledge of all candidate phases is updated (i.e., the present phase fraction at a given processing condition). Therefore, the uncertainty in the phase map is iteratively reduced over all candidate phases simultaneously.

Figure 3(a) presents the final, labeled phase diagram of Bi_2O_3 after 95 iterations. This phase diagram is derived from expected phase fractions, using color saturation to indicate the relative abundance of each phase. A comparison with panel D in Fig. 3 of Ref. [58] shows good agreement between the phase maps. There is a temperature shift compared to previous results due to the improvement in the temperature calibration method (see Sec. IV C). The Bi-O system is particularly challenging for automatic phase labeling algorithms as the XRD patterns of $\beta\text{-Bi}_2\text{O}_3$ are very similar to $\delta\text{-Bi}_2\text{O}_3$ and differ only by some peak splitting due to symmetry breaking. This similarity makes determining the coexistence region through XRD quite difficult. However, the probabilistic labeling helps to identify the two-phase region, leading to creation of a reliable phase diagram.

Note that the crystallographic textured $(\delta + \beta)\text{-Bi}_2\text{O}_3$ observed above approximately 700 °C arises from the melt-quenching process. The pronounced texturing results in the absence of many peaks from the XRD data, rendering the NMF and phase labeling less reliable. Consequently, assistance from a human expert is required to clearly identify such regions. This can be accomplished either on the fly during the AL campaign or, as in our case, as a postprocessing step while analyzing the data to produce the final phase diagram. While dealing with XRD data from such textured films may occasionally produce erroneous labeling within our workflow, we rarely expect significant texturing or the formation of grain sizes comparable to the x-ray probe size, both of which can lead to missing peaks in XRD patterns before the melting temperature is reached. On the other hand, if texturing conditions are expected *a priori*, it is possible to reliably characterize

these features on the fly by incorporating the mediation technique outlined in the original work on CrystalShift [42].

The learning behavior of SARA-H's AL agent is shown in Fig. 3(b), where the R_s^2 (see Sec. IV H 2) score at each iteration is plotted in relation to the fully converged phase diagram constructed from the collected XRD data from 95 anneal experiments. The figure shows that after only some 50 iterations, the average R_s^2 score reaches a value exceeding 0.8, indicating that the phase diagram has already converged. A similar value of R_s^2 was achieved after more than 70 iterations in Ref. [58] for the same materials systems, which is 40% more than what is observed in this work.

The reason for this faster learning can be attributed to two key factors: (1) Accelerated learning is facilitated by the increased richness of physical information contained in the XRD data compared to reflectance spectrometry. Rather than relying solely on the rate of change in the optical data, the XRD provides phase information at different conditions, allowing direct optimization of target phases. (2) The objective of maximizing phase activations is significantly easier to optimize than focusing solely on phase boundaries based on optical gradients, as the phase space occupied by boundary regions constitutes only a fraction of the phase fields themselves. This issue is further exacerbated when moving toward systems with increasing compositional degrees of freedom, as the likelihood of sampling a phase boundary region compared to a phase field diminishes further.

We also applied SARA-H to a second single-composition system, the SnO_x suboxide system, as a further benchmark to reproduce and validate the key findings of prior work by Sutherland [72]. Using the same capped and uncapped suboxidized Sn-O thin films, SARA-H's autonomous learning framework recovered processing phase diagrams that closely match those previously obtained through exhaustive lg-LSA anneals and synchrotron-based experiments. The learned phase behavior reproduces the central conclusions of Sutherland [72], including the sequence of phase transformations from cocrystallized SnO /rutile SnO_2 to full oxidation and eventual decomposition to metallic Sn, as well as the decisive role of atmospheric oxygen uptake during millisecond-scale anneals. In particular, uncapped films rapidly oxidize to SnO_2 , whereas an Al_2O_3 capping layer effectively suppresses oxygen ingress and stabilizes SnO across a broad processing window. The autonomously learned phase diagrams of the capped and uncapped samples are shown in Fig. S3 in the Supplemental Material [43], while the associated learning curves shown in Fig. S4 in the Supplemental Material [43] demonstrate rapid convergence, requiring fewer than 10 iterations for both the uncapped and capped films to reach $R_s^2 = 0.8$, indicating that the SnO_x phase space is comparatively straightforward for SARA-H to learn. Relative to the 105 experiments required in the original exhaustive study, this represents a resource reduction exceeding an order of magnitude. Full experimental details and analyses are provided in Sec. C of the Supplemental Material [43].

2. Active learning on a pseudobinary composition spread: Bi-Ti-O

A much more challenging autonomous experiment involves combining Bi_2O_3 with TiO_2 . To date, more than eight

TiO_2 polymorphs are known, with the most common crystalline phases comprising rutile ($\alpha\text{-TiO}_2$), anatase ($\beta\text{-TiO}_2$), and brookite ($\gamma\text{-TiO}_2$). The latter two are metastable but are preferred in practical applications due to their higher photocatalytic activity [51–53,73,74].

The formation of various compounds in the ternary $\text{Bi}_2\text{O}_3\text{-TiO}_2$ system has been documented in the literature, garnering interest particularly due to the emergence of several material properties for diverse applications, including superior dielectric, piezoelectric, photorefractive, photocatalytic, or ferroelectric characteristics [75–81]. Kargin *et al.* [71] conducted a comprehensive study that summarized results from the literature in the form of an up-to-date equilibrium phase diagram for the $\text{Bi}_2\text{O}_3\text{-TiO}_2$ system, presented in Fig. 4(b) for reference. A full review of this material system is included in Sec. D of the Supplemental Material [43]. Based on the literature research, we set the goals of studying this system to be verifying the prediction of stabilization of anatase with Bi doping and studying the dynamics of binary oxide formation in this system.

We utilized SARA-H to investigate a thin film library of TiO_2 and Bi_2O_3 , which was deposited by rf-reactive sputtering using elemental targets located at geometrically opposed positions. This synthesized a compositional-spread thin film spanning most of the $\text{Bi}_2\text{O}_3\text{-TiO}_2$ pseudobinary, as shown in Fig. S6 of the Supplemental Material [43]. To initiate the AL campaign, experimentally resolved crystal structures from the literature were gathered and 12 candidate phases were selected to seed both CrystalShift and the AL agents. In addition to crystal structures of the unary and binary oxide of Bi-Ti-O systems, SiBi_2O_5 was also included, as Bi_2O_3 had been observed to react with the substrate to form a silicide upon melting. With the exception of $\text{Bi}_{36}\text{Ti}_{32}\text{O}$, all candidate phases lie on the tie line between Bi_2O_3 and TiO_2 in the Bi-Ti-O Gibbs triangle, as shown in Fig. 4(a), ensuring the nominal oxidation state and associated charge balance.

In contrast to unary systems, we progressively adjusted the candidate phase pool to guide the synthesis process for the Bi-Ti-O system, a procedure that roughly divides the AL campaign into three stages (Fig. 5):

(1) Exploration (45 iterations): At the initial stage of the AL campaign, we deployed the cycle sampling acquisition function, considering all 12 candidate phases. The full set of candidate phases is listed in the Supplemental Material [43].

(2) Refined exploration (25 iterations): After initial exploration, the phase fields were visualized and inspected. Some phases were determined to be unsynthesizable using lg-LSA on our amorphous thin-film precursor and were therefore removed to expedite the process.

(3) Targeted exploitation (50 iterations): In the final stage, three phases were retained to focus on the objective that we had initially set.

The phase-resolved learning curves of SARA-H's AL campaign are shown in Fig. S5 of the Supplemental Material [43]. In contrast to the single-composition samples of Bi_2O_3 and SnO_x , the Bi-Ti-O system is considerably more complex and challenging to learn, as evidenced by the much slower increase in the average R_s^2 score. A total of approximately 80 iterations were needed for the average R_s^2 to exceed $R_s^2 = 0.8$. Another subtle detail to note is the sudden increases in the R_s^2




FIG. 4. (a) Gibbs triangle of the ternary phase space Bi-Ti-O. The diagram features all compositions considered in the candidate pool for our probabilistic phase determination algorithm, represented as either circles or crosses. Circles indicate compositions that were observed in the experiment, while crosses denote phases that did not form. Note that most candidate compositions lie along the tie line between Bi_2O_3 and TiO_2 , thereby ensuring appropriate charge balance. (b) Equilibrium phase diagram based on Ref. [71]. The compositional axis is rescaled from the original figure in Ref. [71] to match the atomic fraction convention employed in this article.

scores for some of the individual phases. These jumps may be associated with the cycle sampling strategy as shown in the benchmark section.

Three-dimensional phase fields can be plotted as three-dimensional volume plots (Fig. S8 of the Supplemental Material [43]), but it is challenging to extract quantitative information from 3D plots. The dwell dimension contributes the least variation in phase fraction, so to facilitate further analysis, we reduced the dimensionality by dividing the dwell domain into two halves and projecting out all τ dependence for each half. The two resulting two-dimensional (2D) phase diagrams for short ($300 \mu\text{s} < \tau < 1000 \mu\text{s}$) and long dwell times ($3000 \mu\text{s} < \tau < 10\,000 \mu\text{s}$) are shown as color-coded density plots in Figs. 6(a) and 6(b), respectively, where the expected phase activations are represented by colored regions. The color saturation in each figure represents the probability

densities, indicating the presence of the phases at the relevant processing peak temperatures and compositions.

An immediately and clear feature of each phase map is the wide range of compositions that are associated with each phase, in marked contrast to the very narrow single-phase regions (i.e., line compounds) observed in equilibrium [e.g., $\text{Bi}_4\text{Ti}_3\text{O}_{12}$ in Fig. 4(b)]. This is readily understood by considering the short time during which the atoms have sufficient thermal energy to diffuse, which can both influence the crystal structure that forms and inhibit segregation into two-phase structures.

The results of short dwells illustrated in Fig. 6(a) reveal an extensive amorphous domain with two extrema: (1) for low temperatures and low Bi content, the initially amorphous film does not overcome the energy barriers for nucleation and crystallization and (2) for high temperature and high Bi content, the film undergoes rapid cooling from the melt, thus quenching in the amorphous structure. The existence of this second region is consistent with the trend of the equilibrium phase diagram [Fig. 4(b)], which shows liquid phases in the high-temperature low Bi-content corner.

For intermediate values of Bi fraction and temperature ($0.5 < x < 0.7$ and $900\text{ }^{\circ}\text{C} < T_p < 1100\text{ }^{\circ}\text{C}$), the films are also found to retain the as-deposited amorphous structure. Careful manual examination of the XRD patterns in this region confirms this conclusion. The reasons for failure to crystallize films near equimolar Bi and Ti composition are not entirely clear; however, we expect that several factors contribute to the formation and shape of this amorphous region.

First, the kinetics of crystallization play an important role: For Bi_2O_3 in Sec. II C 1, the crystallization onset at a dwell time of $\tau = 10\,000\text{ }\mu\text{s}$ is $T = 280\text{ }^\circ\text{C}$, whereas the corresponding crystallization onset occurs at $T = 480\text{ }^\circ\text{C}$ for TiO_2 (observed on a single composition TiO_2 thin-film sample). It is thus not surprising that as-deposited Bi_2O_3 -rich films exhibit lower activation energy for the nucleation of crystallites compared to TiO_2 .




FIG. 5. Stages of active learning for the Bi-Ti-O autonomous experimentation.




FIG. 6. The experimentally determined processing phase diagrams of the Bi-Ti-O system after SARA-H's converged AL cycles with 120 iterations. Panel (a) displays the phase activations for the short dwell times of $\tau \in [300 \mu\text{s}, 1000 \mu\text{s}]$. Colors representing different phases are randomly selected from the Hue-Saturation-Value (HSV) hue wheel. The intensity of each region is proportional to the corresponding phase activation. The amorphous domain is represented as the uncolored region. The dashed black lines denote the presumed phase boundaries and are merely intended as a guide for the eye. The cross hatch indicates the model's uncertainty and is more visible when the uncertainty is higher. Panel (b) shows the phase activations for the long dwell times of $\tau \in [3000 \mu\text{s}, 10\,000 \mu\text{s}]$.

Second, we anticipate a reduced driving force from the as-deposited amorphous state to the crystalline phases due to the increased compositional entropy when moving toward the region where both Bi and Ti are present in comparable concentrations. In the compositional range near $x = 0.5$, the entropy of mixing is higher, and the free energy difference to a low-entropy, ordered state is reduced. Consequently, the dwell times of our anneals are too short compared to the required nucleation timescales.

As expected, for the most Bi- and Ti-rich compositions, polymorphs of the unary oxides TiO_2 and Bi_2O_3 are formed. However, the $\delta\text{-Bi}_2\text{O}_3$ phase, which persists only in a narrow temperature range for pure Bi_2O_3 (see also Sec. II C 1), dominates the phase diagram for $x > 0.1$. This is not surprising since substitution by aliovalent ions is well known to stabilize the δ phase [82–86].

Importantly, the pyrochlore phase $\text{Bi}_2\text{Ti}_2\text{O}_7$ occurs over a wide composition range ($0.2 < x < 0.4$) on the Bi-rich side of the $\text{Bi}_4\text{Ti}_3\text{O}_{12}$ phase region, despite having a lower Bi fraction. Evidently, in this composition range, the film first crystallizes into the Ti-substituted cubic $\delta\text{-Bi}_2\text{O}_3$ structure (an inference based on the assumption that the structure observed after quenching from peak temperatures below 900°C is formed during heating). For regions heated to peak temperatures above 900°C , the cubic structure transforms into pyrochlore $\text{Bi}_2\text{Ti}_2\text{O}_7$ even though the thermodynamically stable phase is $\text{Bi}_4\text{Ti}_3\text{O}_{12}$ and the stable composition range for $\text{Bi}_2\text{Ti}_2\text{O}_7$ is very narrow [see Fig. 4(b)]. We speculate that it is kinetically favorable to transform to the metastable pyrochlore structure rather than the Aurivillius phase structure (which has a lower free energy), as the latter requires more atomic hops to form.

For Ti-rich ($x > 0.6$) compositions, we observe an intricate progression of phases as the peak temperature is

increased. For $900^\circ\text{C} < T_p < 1000^\circ\text{C}$, the amorphous film crystallizes into Bi-substituted anatase $\beta\text{-TiO}_2$, a metastable structure that is typically observed at low temperature. At higher temperatures ($1000^\circ\text{C} < T_p < 1100^\circ\text{C}$), a mixture of anatase and rutile structure is obtained, and at the very highest temperatures ($T_p \geq 1300^\circ\text{C}$) only rutile $\alpha\text{-TiO}_2$ remains. Interestingly, at intermediate temperatures, $1100^\circ\text{C} < T_p < 1300^\circ\text{C}$, $\beta\text{-TiO}_2$ is no longer present, but we observe $\alpha\text{-TiO}_2$ coexisting with $\text{Bi}_2\text{Ti}_4\text{O}_{11}$, a phase expected from the equilibrium phase diagram in Fig. 4(b), which nevertheless is not present at lower or higher temperatures.

In comparing the results of the longer dwell time to those of the short dwell times, it is readily apparent that the long-duration synthesis phase diagram [see Fig. 6(b)] exhibits more phases and larger crystalline phase fields. In contrast to the case for short dwells, there are two disconnected amorphous regions: The intermediate amorphous region observed for short dwells defined by $0.5 < x < 0.7$ and $900^\circ\text{C} < T_p < 1100^\circ\text{C}$ in Fig. 6(a), now is part of an expanded phase field for $\text{Bi}_2\text{Ti}_2\text{O}_7$. This is consistent with our interpretation that it is kinetics that inhibits crystallization. For longer dwells, there is sufficient time for atoms to overcome activation barriers to nucleate crystallization.

The most striking feature of the long-dwell-time phase diagram is that the pyrochlore $\text{Bi}_2\text{Ti}_2\text{O}_7$ completely dominates the space within $0.4 < x < 0.8$ for $T_p > 800^\circ\text{C}$. This is unexpected since $\text{Bi}_2\text{Ti}_2\text{O}_7$ appears as a line compound in the equilibrium phase diagram. These results suggest that $\text{Bi}_2\text{Ti}_2\text{O}_7$ nucleates first across a wide range of compositions due to a low activation barrier before it can transform into thermodynamically more stable phases at timescales much longer than those of our experiments.

For the most Ti-rich compositions, we observe a progression similar to that seen at shorter dwells, but the stability

and phase-pure range of β -TiO₂ is increased significantly. Figure S7 of the Supplemental Material [43] illustrates the way in which the relatively small phase volume identified as β -TiO₂ evolves during the AL cycles.

The presence of substantial regions of pure β -TiO₂ in our samples provides strong evidence that Bi acts as a dopant that inhibits the transformation of anatase to rutile in TiO₂. This finding aligns well with the results of Hanaor *et al.* [56], who predict that Bi will function as an inhibitor. To the best of our knowledge, the present work is the first experimental evidence that doping TiO₂ with Bi inhibits the anatase-to-rutile transformation.

Finally, we did not observe any indication of sillenite Bi₁₂TiO₂₀, even after a thorough postanalysis of the XRD patterns from all anneals, neither at short nor at long dwell times. This observation contrasts with the equilibrium phase diagram in Fig. 4(b) and reports by Speranskaya *et al.* [87] and Maier *et al.* [88] (as reviewed in Kargin *et al.* [71]). This finding is not entirely surprising, as the compositional domain where Bi₁₂TiO₂₀ should form is dominated by δ -Bi₂O₃, which, as discussed, readily forms as the primary nucleating phase and which is capable of accommodating a significant level of Ti impurities due to its defective fluorite structure.

Overall, active learning of the synthesis phase diagram in this pseudobinary system has revealed a variety of unanticipated and complex phase evolution behaviors. Rather than following simple or monotonic trends, the observed phase behavior reflects the intricate interplay between synthesis conditions, kinetic pathways, and metastable phase formation. Extending this approach to accumulate larger and more diverse datasets across a broader range of chemical systems will enable the identification of recurring patterns and correlations. Ultimately, such systematic, data-driven exploration is expected to facilitate the development of a more unified and predictive understanding of the principles governing metastable phase synthesis.

III. CONCLUSION

We have significantly enhanced the capabilities of SARA, our AI-driven autonomous closed-loop materials discovery framework, by integrating on-the-fly XRD analysis and structural characterization with robotic materials synthesis using lg-LSA. By leveraging an efficient and robust algorithm for probabilistic phase labeling, as implemented with NMF and CrystalShift, the enhanced SARA-H can now not only classify the crystal structures of synthesized metastable phases in a timescale commensurate with the pace of laser annealing and XRD experiments, but also autonomously conduct a targeted search for desired phases with high purity. SARA-H's capabilities were highlighted by investigating three material systems, namely, Bi₂O₃, SnO_x, and Bi-Ti-O. For Bi-Ti-O, our targeted search notably resulted in identifying the processing space to form anatase, a polymorph of TiO₂ that is particularly sought after due to its desirable photocatalytic properties and is stabilized by Bi doping.

To achieve efficient learning, we developed a novel sampling strategy based on the logEI acquisition function that balances the trade-off between the exploration and exploitation of aggregated information during AL experimentation,

which we refer to as cycle sampling. Extensive benchmarks demonstrate that cycle sampling significantly outperforms the baseline sampling methods (random and uncertainty sampling). Most importantly, its target objective can be readily modified during an experimental campaign (i.e., on the fly) through the intervention of a human expert who can steer the search toward the targeted synthesis of specific phase(s). Careful integration of a human in the loop in AL designs not only accelerates the time to solution but has proven to be crucial in other AI domains where the well-orchestrated interplay between human and machine can exploit their individual advantages to ultimately optimize complex discoveries [38,39].

While this work focuses on optimizing the targeted AL cycles specifically to maximize desired phase(s), the objective property can easily be substituted. For instance, we might aim to maximize unclassifiable diffraction patterns for the discovery of novel materials. Alternatively, we could incorporate additional characterization agents for property measurements, such as optical spectrometry to evaluate absorption for identifying synthesized materials suitable for solar energy applications, Raman spectroscopy that would reveal structural and chemical properties not exposed by XRD, or microscopic four-point probes for resistivity measurements. A wide range of processing-structure-property relationships can be revealed by integrating these agents into SARA-H's AL cycles, and developing effective methods is a promising area for AI research.

Another avenue for further research involves integrating emerging LLM agents directly into the overarching SARA workflow. Utilizing these models as dynamic knowledge bases offers a novel pathway for uncovering latent conceptual links and surfacing hidden hypotheses previously overlooked in traditional literature reviews. Ultimately, this approach aims to augment, rather than replace, human reasoning. Several scientific workflows have already started to incorporate LLM into their closed-loop experimental systems [89,90]. However, deploying LLMs as cognitive accelerators will require robust experimental guardrails and standardized benchmarks to rigorously evaluate AI-generated insights against expert human analysis.

The hierarchical multiagent system described here integrates experimental materials processing and characterization with AI-based algorithms for reasoning and scientific discovery, including the representation, planning, optimization, and learning of materials knowledge. Our experimental methods are conducive to the formation of metastable materials which, broadly speaking, can exhibit unique and technologically valuable properties. However, the synthesis pathways, stability ranges, and structure-property relationships of metastable phases remain poorly characterized, largely due to the difficulty of systematically exploring the vast and dynamic processing space in which they form. By autonomously generating, characterizing, incorporating human-in-the-loop feedback, and iteratively refining experimental conditions, our system can produce rich, high-dimensional datasets that capture the emergence and evolution of metastable states. With its capabilities, SARA-H acts as a powerful catalyst for the discovery of metastable materials, bridging the gap between theoretical potential and the rapid realization of next-generation technologies.

IV. METHODS

A. Thin-film libraries

Composition libraries were deposited on heavily *p*-type doped ($0.01 \Omega \text{ cm}$) silicon wafers ($500 \mu\text{m}$ thick). Unless otherwise noted, all substrates had a 20 nm thermally grown oxide layer. Gold alignment features were first formed on wafers using standard lithography. Thin films were then deposited on the wafers using rf magnetron sputtering using 2-in. diameter targets for the corresponding material system. The atmosphere and detailed sputtering conditions for each film are as follows:

(1) *Bi-O system*: rf sputtering from a Bi target in an atmosphere of 8 mTorr Ar and 2 mTorr O_2 to deposit the Bi_2O_3 sample on 200 nm thermally oxidized SiO_2 wafers. Wafers were rotated while sputtering at an rf power of 20 W to produce a 170 nm thick film with $<10\%$ thickness variation across the wafer.

(2) *Sn-O system*: Thin films of SnO_x were sputtered from a metallic Sn target in an O_2 and Ar gas mixture of $6\% \text{ O}_2$ at a total pressure of 5 mTorr in an AJA International deposition system. Samples were rotated to form $\approx 100 \text{ nm}$ thick, predominantly amorphous, films. Half of the wafer was capped with a $\approx 15 \text{ nm}$ thin film of Al_2O_3 sputtered from an Al_2O_3 target to reduce O_2 exposure of the film during laser annealings.

(3) *Bi-Ti-O system*: The $(\text{Ti}, \text{Bi})\text{O}_x$ composition spread films were sputtered using 2-in. diameter metal targets in two Angstrom [91] sputter guns arranged in the 90° off-axis configuration to minimize bombardment of the growing film by energetic O anions in the plasma [3]. The Ti target ($>99.9\%$ pure) was rf powered at 125 W and the Bi target ($>99.99\%$ pure) at 15 W rf, while the Si substrate was rf powered at 6 W , resulting in a dc bias of -40 V to help densify the as-grown film. The sputtering ambient was $20\% \text{ O}_2$, $80\% \text{ Ar}$ supplied at an overall flow rate of 30 SCCM at 30 mTorr . The substrate was unheated during sputtering and as-deposited films were amorphous. The cation composition range was roughly $0.1 < \text{Bi}/(\text{Bi} + \text{Ti}) < 0.9$ (detailed compositional characterization postdeposition is reported in Sec. IV B together with the XRF measurements in Sec. D of the Supplemental Material [43]). The film was deposited for 60 min , yielding thicknesses ranging from 200 to 300 nm over the composition spread. As with Ti-O, no thickness-dependent variations in phase formation are observed within this range.

B. Compositional characterization for Bi-Ti-O

The composition of binary spreads was evaluated using XRF using an EDAX Orbis Micro-XRF system with a 2 mm x-ray beam diameter. The Bi K and Ti K XRF peak intensities were converted to normalized cation compositions using sensitivity factors calibrated by commercial XRF standards (MicromatterTM). The measured composition of the film is presented as a colored scatter plot in Fig. S6 of the Supplemental Material [43]. Since the density of the XRF data was lower than that of the Ig-LSA stripes on the wafer, the data were interpolated using a smooth GP model with prior length scales of $\sigma_{x,y} = 2 \text{ mm}$ and an RBF kernel.

C. Lateral-gradient laser spike annealing

The Ig-LSA was based on the design introduced by Bell *et al.* [41]. A 120 W CO_2 laser ($\lambda = 10.6 \mu\text{m}$) was passed through an attenuator to control the incident power on the wafer, calibrated with a power meter positioned at the wafer plane. The incidence angle of the laser with respect to the wafer normal was 73° , close to Brewster's angle to minimize reflected power. Due to this shallow incidence angle, the temperature profile exhibited a two-sided Lorentzian intensity distribution on the wafer, with the laser intensity displaying different widths on either side of the peak power.

To secure the wafer in place and ensure consistent thermal contact, a porous aluminum vacuum wafer chuck was used. For each anneal, the stage was first accelerated to the desired velocity corresponding to the desired dwell time. Once the velocity was reached, a shutter opened to anneal the sample. After the anneal stripe reached the designed length (5 mm), the shutter was closed, and the stage was decelerated to a complete stop.

D. Temperature calibration of Ig-LSA

To extract the dense temperature-phase information from Ig-LSA stripes, we established an accurate calibration of the full temperature profile generated during annealing. For a given laser power P and scan velocity v , the goal is to determine the peak anneal temperature distribution $T_p(v, P, x)$ at every position x across the stripe. We calibrated this temperature profile at a fixed downstream position y along the stripe, namely, at the $y = 4 \text{ mm}$ mark of an anneal stripe, chosen to ensure quasiequilibrium laser conditions while avoiding artifacts associated with shutter actuation.

Accurate temperature calibrations were obtained by employing a thermorefectance approach, which leverages the linear dependence of silicon's optical reflectivity on absolute temperature change [93]. By monitoring relative reflectance changes during laser scans with a camera and associating them with a reference reflectance at a known temperature, we obtained a robust, spatially resolved measure of the peak thermal response under each annealing condition.

A comprehensive calibration dataset was assembled by repeating the reflectance measurements over a wide range of laser powers and scan velocities and was subsequently used to fit a parametrized temperature model. Instead of modeling the peak anneal temperature distribution $T_p(v, P, x)$ as a single function, we constructed empirical functional relationships that describe the peak temperature at the center of the stripe ($T_c(v, P)$) and the lateral extent along x of the temperature distribution, modeled as a two-sided Lorentzian curve with widths $R(v, P)$ and $L(v, P)$ for the right and left sides, respectively. To convert relative temperatures from reflectance measurements into an absolute temperature scale, the model was calibrated by experimentally determined melting conditions of silicon and gold, whose melting points are well established. Once calibrated, the center peak temperature and the width distributions were converted from scan velocities v to dwell times $\tau = \frac{\text{FWHM}}{v}$, where FWHM denotes the full width at half maximum of the laser power profile along the scanning direction y . The combined hypersurfaces of $T_c(\tau, P)$, $R(\tau, P)$, $L(\tau, P)$ enable prediction of complete temperature

profiles as a function of laser power and dwell time with an accuracy of approximately ± 25 °C near the center of a stripe. A detailed description of thermoreflectance modeling, fitting procedures, and calibration strategy was described elsewhere [92].

E. High-throughput wide angle x-ray diffraction

The lg-LSA anneal stripes are nominally ≈ 1.5 mm wide, resulting in temperature gradients of $\approx 1\text{--}5$ °C/ μm , depending on laser anneal conditions. To sample this temperature distribution with sufficient spatial resolution and in a high-throughput manner, we utilized wide-angle x-ray scattering at the Cornell High Energy Synchrotron Source (CHESS) ID3B beamline [94], employing a 9.7 keV x-ray beam focused by a compound reflective lens to a projected spot size of 20×60 μm upon striking the thin-film sample at a 3° incident angle. The diffracted x-ray signals were collected by a Dectris Eiger 1M detector.

To collect spatially resolved x-ray diffraction across an annealed stripe, the detector continuously gathered 2D x-ray diffraction data at 50 ms intervals while the lg-LSA precision stage stepped horizontally relative to the x-ray beam location. The stepping velocity was set to achieve a spatial resolution of 10 μm , which corresponds to a temperature resolution better than $\Delta T \approx 50$ °C. A total of 151 two-dimensional x-ray diffraction images were collected for each anneal stripe, covering a distance of 1.5 mm. An open-source azimuthal integration package, PyFAI [95], was used to integrate the 151 2D x-ray diffraction images in parallel into 1D XRD patterns on the fly, utilizing the high-performance computing facility at CHESS. It is critical to carefully characterize the scattering geometry since small uncompensated variations can lead to large errors in the scattering vector due to the low angle of incidence of the x-ray beam. A representative heat map of a scanned XRD map is shown in the lower portion of the second panel of Fig. 1.

F. Automated x-ray diffraction analysis and phase labeling

In the first step of processing the acquired, integrated XRD patterns of an annealed stripe, NMF was performed to reduce the dimensionality of the data and the subsequent demand on computational resources. The NMF can be accomplished by solving the following approximation problem:

$$\begin{aligned} \mathbf{X} &\approx \mathbf{WH}, \\ \text{s.t. } \mathbf{W} &\geq 0, \quad \mathbf{H} \geq 0, \end{aligned} \quad (1)$$

where $\mathbf{X} \in \mathbb{R}^{n \times m}$ is the data matrix, in which each column represents an XRD pattern, and the inequalities apply elementwise. For this reason, $\mathbf{W} \in \mathbb{R}^{n \times k}$ is a matrix whose columns can be interpreted as prototypical non-negative XRD patterns that are added together in different proportions using the non-negative coefficients in $\mathbf{H} \in \mathbb{R}^{k \times m}$. Due to the matrices' functions, we refer to $\mathbf{W}$ as the basis matrix and $\mathbf{H}$ as the activation matrix. The integer k is a user-defined parameter that represents the rank of the data matrix. In the present implementation, we applied an NMF that is based on an extreme ray finding algorithm [60]. Notably, the objective of this algorithm is to select representative columns

in $\mathbf{X}$ to construct a basis set that minimizes Eq. (1), rather than to separate the phases in a mixture. We addressed this separation problem with other methods below. This selection process avoids the formation of nonphysical basis XRD patterns. In all our AL campaigns, we used $k = 4$ and attributed the basis that has the highest activation value at the first column of the XRD data to be the signal from the as-deposited film and discarded it from subsequent data processing.

To reduce computational cost and avoid labeling noncrystalline signals, we fit each basis with a smooth background model using kernel ridge regression (RBF kernel). Bases with fitting errors below a threshold were classified as amorphous and discarded. The threshold was set conservatively to ensure that crystalline signals are not inadvertently removed.

Next, CrystalShift, a rapid probabilistic multiphase labeling algorithm was employed to label phases and estimate the phase fractions associated with each basis vector in $\mathbf{W}$ [42]. CrystalShift effectively integrates a best-first tree search, rapid lattice refinement, and Bayesian model comparison to provide probability estimates for potential phase combination assignments.

Here, we briefly describe the high-level concept and process of the algorithm, with further details available in the original papers. CrystalShift uses a best-first tree search scheme to explore the combinatorial space of phases. In the tree, each node represents a possible phase combination. At each level of the tree, CrystalShift performs a pseudorefinement process on all nodes at that level. For each factorized XRD pattern basis vector $\mathbf{W}_i$, the pseudorefinement algorithm minimizes the difference between the simulated diffraction pattern of the phase combination and the given pattern by modifying the crystal system-preserving parameters θ of a given candidate phase set

$$\theta^* = \arg \min_{\theta} \|f(\theta) - \mathbf{W}_i\|_2^2, \quad (2)$$

where f is the function that computes the XRD pattern from the parameters θ . Then, the nodes with the k_{exp} lowest residuals are expanded, which means adding an additional phase that was not originally present in the node to the next level. In this work, $k_{\text{exp}} = 3$ was used. Starting with an empty root node at level 0, CrystalShift repeats the refine-and-expand loop until the desired level is reached, as defined by the maximum number of allowed coexisting phases specified by the user.

After the tree search is completed, the likelihood of each node is calculated by marginalizing out the parameters θ of the model. Since this process is mathematically intractable, the Laplace approximation is used to perform the marginalization. From Eq. (2), the log likelihood of a set of candidate phases M_i can be written as

$$\ln p(\mathbf{t}|M_i, \theta_i^*) = -L_i, \quad (3)$$

where $L_i = \|f(\theta_i^*) - \mathbf{W}_i\|_2^2$ is the residual of Eq. (2). Using Laplace approximation to marginalize out θ gives

$$\ln p(\mathbf{t}|M_i) \approx -L_i - \frac{\ln |H_{\theta_i^*}(L_i)|}{2} + d_i \frac{\ln(2\pi)}{2}, \quad (4)$$

where $H_{\theta_i}^*(L_i)$ is the Hessian of the loss with respect to the variables at the optimum and d_i is the dimensionality. To ensure the probabilistic labeling was well calibrated, a temperature scaling was used [96], which works by multiplying the resulting log likelihood by a temperature parameter β . The temperature parameter was determined by performing thousands of labelings on a calibration dataset containing 10 000 synthetic XRD patterns that approximately span the same q range as our experiment. The XRD calibration dataset was generated using candidate phases used in each AL trial and each pattern contained up to two phases with symmetry-allowed lattice strains $\epsilon \in [-0.05, 0.05]$, phase fractions $p_i \in [0.25, 0.75]$, and a random peak width factor $\sigma \in [0.1, 0.3]$. The temperature parameter was selected such that it minimizes the mean absolute error between the calibration curve and the ideal diagonal line in the reliability diagram [97]. Finally, the scaled log likelihood passes through a softmax function to produce a probability distribution over all possible phase combinations given the experimental XRD pattern. The posterior probability for each set of candidate phases M_i is then given by

$$p(M_i|\mathbf{t}) = \frac{\exp(\beta \ln p(\mathbf{t}|M_i))p(M_i)}{\sum_{j=1}^N \exp(\beta \ln p(\mathbf{t}|M_j))p(M_j)}, \quad (5)$$

where usually the prior over phase combinations is uniform, i.e., $p(M_i) = 1/N$.

To enhance SARA-H's robustness against experimental uncertainty and phase labeling errors, we provide expected phase activations that incorporate the probability estimates generated by CrystalShift, rather than directly reporting the most probable phase fraction to the AL agent. A combinatorial sum over the top candidates of each basis was performed to combine the results. Here, we considered only the top 5 candidates for each of the 3 bases and produce 125 potential phase models for the stripe. The final phase probabilities produced by this model are then the product of all probability estimates of each phase model M_i , renormalized so that the probabilities of all 125 models sum to 1.

Finally, we calculate the expected phase activations of all phases by

$$\mathbb{E}[a_j^{(i)}] = \sum_{k=1}^K H_{kj} \mathbf{p}^T \mathbf{a}^{(i)}, \quad (6)$$

where $a_j^{(i)}$ represents the activation of phase i at position index j , H_{kj} is the activation matrix element at the k th row and j th column, $\mathbf{p}$ is the probability vector of the 125 models, and $\mathbf{a}^{(i)}$ is the collected vector of the activation of phase i in each of the 125 models. The expected activation of phase i at location index j , $\mathbb{E}[a_j^{(i)}]$, was then mapped to the expected activation of phase i at temperature T , $\mathbb{E}[a_j^{(T)}]$, using the temperature profile introduced in Sec. IV D. The activation was smoothed using a GP with a square exponential kernel. To avoid secondary effects that may occur at the edges of the XRD map, i.e., the tails of the temperature profile, and to ensure that the temperature gradient is not too steep and, therefore, that the thermal resolution was sufficiently high, we sampled 16 points between the peak temperature and 200 °C below the peak temperature with equal spacing.

G. Active learning for targeted phase mapping

1. Gaussian process models

The goal of active learning is to reason with the currently available expected phase activation distribution to suggest the next experimental condition that would yield the largest expected improvement of the phase diagram. We utilized an AL scheme based on Bayesian optimization, which includes a surrogate model and an acquisition function for experimental decision making. In our cyclic AL setup, a GP regression model was chosen as the surrogate model because of its reliable uncertainty estimates and its capability to incorporate prior domain knowledge through kernel function constructions.

A GP is a distribution over functions, and its behavior is entirely determined by its mean $\mu(\cdot)$ and covariance function $K(\cdot, \cdot)$, commonly referred to as the kernel function. Given a set of currently available input data in matrix form $\mathbf{X}$, a GP can be expressed as [98]

$$f_{\mathcal{GP}}(\mathbf{X}) \sim \mathcal{N}(\mu(\mathbf{X}), K(\mathbf{X}, \mathbf{X})), \quad (7)$$

where $K(\mathbf{X}, \mathbf{X}')$ is the covariance matrix, with its $(i$ th, j th) entry equal to $K(\mathbf{X}_i, \mathbf{X}_j)$. When performing GP inference on a new point $\mathbf{x}^*$, the GP posterior distribution is generated by conditioning on the joint distribution of the GP prior to obtain the posterior. This process results in another multivariate Gaussian distribution f_p ,

$$f_p = p(f|\mathbf{X}, \mathbf{x}^*, \mathbf{y}, f_{\mathcal{GP}}) = \mathcal{N}(\mu_p(\mathbf{x}^*), \Sigma_p(\mathbf{x}^*, \mathbf{x}^*)), \quad (8)$$

$$\mu_p(\mathbf{x}^*) = K(\mathbf{x}^*, \mathbf{X})K(\mathbf{X}, \mathbf{X})^{-1}(\mathbf{y} - \mu(\mathbf{X})), \quad (9)$$

$$\Sigma_p(\mathbf{x}^*, \mathbf{x}^*) = K(\mathbf{x}^*, \mathbf{x}^*) - K(\mathbf{x}^*, \mathbf{X})K(\mathbf{X}, \mathbf{X})^{-1}K(\mathbf{X}, \mathbf{x}^*). \quad (10)$$

The kernel used in this work was a Matérn-5/2 kernel with length scales of $\sigma_{T_p} = 50$, $\sigma_{\tau_{\log}} = 0.3$, and $\sigma_c = 0.15$ for the dimensions of peak temperature T_p , the decadic logarithm of the dwell time τ , and the composition axis c , if applicable, respectively. In principle, the length scales of the Matérn kernel can be optimized by maximizing the marginal likelihood [98]. However, our selection of the aforementioned length scales was based on benchmarks utilizing a range of fixed values, from which we chose the best-performing combination for the kernel [58].

2. Uncertainties

It is critical to consider all sources of uncertainty when designing nested learning strategies that rely on experimental, noisy data. As outlined in our previous work [58], we accounted for errors in the measurements (i.e., model outputs from the XRD characterization) as well as intrinsic uncertainties in the temperature profile. We approximated the posterior inference with input noise by propagating the input uncertainty using a linear approximation of the standard posterior mean, following the approach of McHutchon and Rasmussen [99]. In this model, given the regular posterior mean $\mu_p(x)$, its input-noise-corrected version can be computed by updating

$$K \leftarrow K + \text{diag}(\sigma_x(\mathbf{X}) \odot \partial_x \mu_p(\mathbf{X}))^2. \quad (11)$$

T_c to $T_c - 200$ °C when probing a condition of (T_c, τ, c) and the acquisition function sums the expected improvement over the same temperature range in order to make use of the temperature gradient for improved decision making. The stripe extension is applied to enhance each of the aforementioned AL methods.

For the benchmark runs, we included 12 phases in the candidate pool. These candidate phases are identical to those used in our real experiments described in Sec. II C 2 (also listed in the Supplemental Material [43]). The Bayesian optimization setup for the AL model was identical to the one reported in Sec. IV G 1, utilizing a Matérn-5/2 kernel with length scales of $\sigma_{T_c} = 50$, $\sigma_{\tau_{\log}} = 0.3$, and $\sigma_c = 0.15$ for the dimensions of peak temperature T_c , the decadic logarithm of the dwell time τ , and the composition axis c , respectively. Each approach was repeated 500 times to gather statistically meaningful averages. The stripe sampling variants were repeated 100 times due to their longer run times.

2. Performance metrics

We evaluated the effectiveness of the active learning strategies using various performance metrics, which are formally defined in this section to provide a clear basis for our benchmarking results. First, the regret r_k for each phase k is defined as $r_k = |a_k(\hat{\mathbf{x}}) - a_k(\mathbf{x}^*)|$, where a_k represents the regression model and maps the parameters within the space of $\{c, \tau, T\}$ to its predicted activation $a_k(\mathbf{x})$. This measure represents the relative difference between the maximum phase activation by the AL model, $a_k(\mathbf{x}^*)$, and its true maximum phase activation, $a_k(\hat{\mathbf{x}})$.

Next, we introduce the IQR of regret, which is defined as the difference between the 75th and 25th percentiles of the regret distribution in the statistical ensemble.

The enhancement factor (EF) at iteration i for phase k is defined by

$$\text{EF}_i^k = \frac{a_{\text{AL},k}(\mathbf{x}_{1:i}^*)}{a_{\text{random},k}(\mathbf{x}_{1:i}^*)}, \quad (15)$$

where $a_{\text{AL},k}(\mathbf{x}_{1:i}^*)$ is the maximum phase activation observed before iteration i of phase k . The EF curve illustrates the ratios between the objective values in the AL runs and the random sampling runs at various stages, indicating how much improvement the AL techniques provide in terms of objective values.

The acceleration factor (AF) is defined as

$$\text{AF}^k(a) = \frac{i_{\text{random},k}(a)}{i_{\text{AL},k}(a)}, \quad (16)$$

where $i_{s,k}(m)$ is given by $\arg \min_i a_{s,k}(\mathbf{x}_{1:i}^*) \geq m$ for phase k with s being the sampling strategy. In other words, the AF measures the speed-up achieved by the AL strategy relative to random sampling in reaching the same objective value.

The coefficient of determination R_k^2 for a phase k is evaluated on the full model, including all of the n acquired data points i , and is defined as

$$R_k^2 = 1 - \frac{\sum_{i=1}^n (a_k(\mathbf{x}_i) - y_i^k)^2}{\sum_{i=1}^n (\mu(\mathbf{y}^k) - y_i^k)^2}, \quad (17)$$

where y_i^k represents the ground truth activation of phase k at coordinate $\mathbf{x}_i$ and $\mu(\mathbf{y}^k)$ denotes the mean of the data $\mathbf{y}^k$. Furthermore, the mean coefficient of determination $R^2 = \bar{R}_k^2$ is the average over R_k^2 of all phases k . However, since we are dealing with heteroscedastic data, we generalized R_k^2 by weighting over the standard deviation σ_i at each data point i , the concept of which is based on the log likelihood of the heteroscedastic normal errors. This approach leads to

$$R_{S_k}^2 = 1 - \frac{\sum_{i=1}^n (a_k(\mathbf{x}_i) - y_i^k)^2 / \sigma_i^2}{\sum_{i=1}^n (\mu(\mathbf{y}^k) - y_i^k)^2 / \sigma_i^2} \quad (18)$$

and its associated mean coefficient of determination $R_S^2 = \bar{R}_{S_k}^2$ as the average over all phases k .

ACKNOWLEDGMENTS

The authors acknowledge the Air Force Office of Scientific Research for its support under Award No. FA9550-18-1-0136. This work was also supported by the Air Force Office of Scientific Research (AFOSR), Defense University Research Instrumentation Program (DURIP) FA9550-23-1-0569, and the Schmidt Sciences AI2050 Senior Fellowship. This work is based on research conducted at the Materials Solutions Network at CHESS (MSN-C), which is supported by the Air Force Research Laboratory under Award No. FA8650-19-2-5220 and the National Science Foundation Expeditions under Award No. CCF-1522054. This work was partially performed at the Cornell NanoScale Facility, a member of the National Nanotechnology Coordinated Infrastructure (NNCI), which is supported by the National Science Foundation (Grant No. NNCI-2025233). This research was conducted with support from the Cornell University Center for Advanced Computing. The authors would also like to acknowledge the technical assistance provided by staff at CHESS.

M.-C.C. and M.A. took the lead in writing the article. M.-C.C., M.A., and S.A. developed and implemented the algorithms to execute SARA-H. D.R.S., R.B.v.D., and K.R.G. fabricated the thin-film samples. L.Z. and J.M.G. provided compositional characterization of the thin films. L.M.S. and A.R.W. offered methodological and technical support for the XRD measurements. R.B.v.D., C.P.G., M.O.T., and J.M.G. supervised the research. M.-C.C. and M.A. contributed equally to this work. All authors provided critical feedback and helped shape the research, analysis, and article.

DATA AVAILABILITY

The data that support the findings of this study are openly available at [104].

[1] C. A. Mirkin, E. H. Sargent, and D. P. Schrag, Energy transition needs new materials, *Science* **384**, 713 (2024).

[2] T. M. I. Mahlia, T. J. Saktisahdan, A. Jannifar, M. H. Hasan, and H. S. C. Matseelar, A review of available methods and

- development on energy storage; technology update, *Renew. Sustain. Energy Rev.* **33**, 532 (2014).
- [3] R. B. v. Dover and L. F. Schneemeyer, The codeposited composition spread approach to high-throughput discovery/exploration of inorganic materials, *Macromol. Rapid Commun.* **25**, 150 (2004).
- [4] J. R. Hattrick-Simpers, J. M. Gregoire, and A. G. Kusne, Perspective: Composition-structure-property mapping in high-throughput experiments: Turning data into knowledge, *APL Mater.* **4**, 053211 (2016).
- [5] S. Kirklin, J. E. Saal, B. Meredig, A. Thompson, J. W. Doak, M. Aykol, S. Rühl, and C. Wolverton, The Open Quantum Materials Database (OQMD): Assessing the accuracy of DFT formation energies, *npj Comput. Mater.* **1**, 15010 (2015).
- [6] A. Jain, S. P. Ong, G. Hautier, W. Chen, W. D. Richards, S. Dacek, S. Cholia, D. Gunter, D. Skinner, G. Ceder, and K. A. Persson, Commentary: The Materials Project: A materials genome approach to accelerating materials innovation, *APL Mater.* **1**, 011002 (2013).
- [7] K. Kim, L. Ward, J. He, A. Krishna, A. Agrawal, and C. Wolverton, Machine-learning-accelerated high-throughput materials screening: Discovery of novel quaternary Heusler compounds, *Phys. Rev. Mater.* **2**, 123801 (2018).
- [8] J. Shen, S. D. Griesemer, A. Gopakumar, B. Baldassarri, J. E. Saal, M. Aykol, V. I. Hegde, and C. Wolverton, Reflections on one million compounds in the open quantum materials database (OQMD), *J. Phys. Mater.* **5**, 031001 (2022).
- [9] A. R. Oganov, *Modern Methods of Crystal Structure Prediction*, 1st ed. (Wiley-VCH Verlag GmbH & Co. KGaA, Weinheim, Germany, 2010).
- [10] C. J. Pickard and R. J. Needs, *Ab initio* random structure searching, *J. Phys.: Condens. Matter* **23**, 053201 (2011).
- [11] M. Amsler and S. Goedecker, Crystal structure prediction using the minima hopping method, *J. Chem. Phys.* **133**, 224104 (2010).
- [12] D. Zagorac, H. Müller, S. Ruehl, J. Zagorac, and S. Rehme, Recent developments in the inorganic crystal structure database: Theoretical crystal structure data and related features, *J. Appl. Crystallogr.* **52**, 918 (2019).
- [13] C. W. Andersen, R. Armiento, E. Blokhin, G. J. Conduit, S. Dwaraknath, M. L. Evans, A. Fekete, A. Gopakumar, S. Gražulis, A. Merkys, *et al.*, OPTIMADE, an API for exchanging materials data, *Sci. Data* **8**, 217 (2021).
- [14] A. W. Senior, R. Evans, J. Jumper, J. Kirkpatrick, L. Sifre, T. Green, C. Qin, A. Židek, A. W. Nelson, A. Bridgland, *et al.*, Improved protein structure prediction using potentials from deep learning, *Nature (London)* **577**, 706 (2020).
- [15] A. Merchant, S. Batzner, S. S. Schoenholz, M. Aykol, G. Cheon, and E. D. Cubuk, Scaling deep learning for materials discovery, *Nature (London)* **624**, 80 (2023).
- [16] K. Choudhary, B. DeCost, C. Chen, A. Jain, F. Tavazza, R. Cohn, C. W. Park, A. Choudhary, A. Agrawal, S. J. Billinge, *et al.*, Recent advances and applications of deep learning methods in materials science, *npj Comput. Mater.* **8**, 59 (2022).
- [17] S. Takeda, A. Kishimoto, L. Hamada, D. Nakano, and J. R. Smith, Foundation model for material science, in *Proceedings of the 37th AAAI Conference on Artificial Intelligence and 35th Conference on Innovative Applications of Artificial Intelligence and 13th Symposium on Educational Advances in Artificial Intelligence*, AAAI'23/IAAI'23/EAAI'23 (AAAI Press, Washington, DC, USA, 2023), Vol. 37, pp. 15376–15383.
- [18] I. Batatia, P. Benner, Y. Chiang, A. M. Elena, D. P. Kovács, J. Riebesell, X. R. Advincula, M. Asta, M. Avaylon, W. J. Baldwin, *et al.*, A foundation model for atomistic materials chemistry, *J. Chem. Phys.* **163**, 184110 (2025).
- [19] B. Settles, Active learning literature survey, Computer Sciences Technical Report 1648, University of Wisconsin–Madison Department of Computer Sciences, 2009.
- [20] A. G. Kusne, H. Yu, C. Wu, H. Zhang, J. Hattrick-Simpers, B. DeCost, S. Sarker, C. Oses, C. Toher, S. Curtarolo, A. V. Davydov, R. Agarwal, L. A. Bendersky, M. Li, A. Mehta, and I. Takeuchi, On-the-fly closed-loop materials discovery via Bayesian active learning, *Nat. Commun.* **11**, 5966 (2020).
- [21] D. Xue, P. V. Balachandran, J. Hogden, J. Theiler, D. Xue, and T. Lookman, Accelerated search for materials with targeted properties by adaptive design, *Nat. Commun.* **7**, 11241 (2016).
- [22] P. Nikolaev, D. Hooper, N. Perea-López, M. Terrones, and B. Maruyama, Discovery of wall-selective carbon nanotube growth conditions via automated experimentation, *ACS Nano* **8**, 10214 (2014).
- [23] P. Nikolaev, D. Hooper, F. Webber, R. Rao, K. Decker, M. Krein, J. Poleski, R. Barto, and B. Maruyama, Autonomy in materials research: A case study in carbon nanotube growth, *npj Comput. Mater.* **2**, 16031 (2016).
- [24] J. Li, S. G. Ballmer, E. P. Gillis, S. Fujii, M. J. Schmidt, A. M. E. Palazzolo, J. W. Lehmann, G. F. Morehouse, and M. D. Burke, Synthesis of many different types of organic small molecules using one automated process, *Science* **347**, 1221 (2015).
- [25] P. J. Kitson, G. Marie, J.-P. Francoia, S. S. Zaleskiy, R. C. Sigerson, J. S. Mathieson, and L. Cronin, Digitization of multistep organic synthesis in reactionware for on-demand pharmaceuticals, *Science* **359**, 314 (2018).
- [26] M. Christensen, L. P. E. Yunker, F. Adediji, F. Häse, L. M. Roch, T. Gensch, G. dos Passos Gomes, T. Zepel, M. S. Sigman, A. Aspuru-Guzik, and J. E. Hein, Data-science driven autonomous process optimization, *Commun. Chem.* **4**, 112 (2021).
- [27] J. S. Manzano, W. Hou, S. S. Zaleskiy, P. Frei, H. Wang, P. J. Kitson, and L. Cronin, An autonomous portable platform for universal chemical synthesis, *Nat. Chem.* **14**, 1311 (2022).
- [28] F. Strieth-Kalthoff, H. Hao, V. Rathore, J. Derasp, T. Gaudin, N. H. Angello, M. Seifrid, E. Trushina, M. Guy, J. Liu, *et al.*, Delocalized, asynchronous, closed-loop discovery of organic laser emitters, *Science* **384**, eadk9227 (2023).
- [29] N. J. Szymanski, B. Rendy, Y. Fei, R. E. Kumar, T. He, D. Milsted, M. J. McDermott, M. Gallant, E. D. Cubuk, A. Merchant, H. Kim, A. Jain, C. J. Bartel, K. Persson, Y. Zeng, and G. Ceder, An autonomous laboratory for the accelerated synthesis of novel materials, *Nature (London)* **624**, 86 (2023).
- [30] A. Wang, A. Aspuru-Guzik, and G. Ozin, Challenges and opportunities for AI in synthetic solid-state inorganic chemistry, *Matter* **7**, 5 (2024).
- [31] J. Leeman, Y. Liu, J. Stiles, S. B. Lee, P. Bhatt, L. M. Schoop, and R. G. Palgrave, Challenges in high-throughput inorganic materials prediction and autonomous synthesis, *PRX Energy* **3**, 011002 (2024).

- [32] D. P. Tabor, L. M. Roch, S. K. Saikin, C. Kreisbeck, D. Sheberla, J. H. Montoya, S. Dwaraknath, M. Aykol, C. Ortiz, H. Tribukait, C. Amador-Bedolla, C. J. Brabec, B. Maruyama, K. A. Persson, and A. Aspuru-Guzik, Accelerating the discovery of materials for clean energy in the era of smart automation, *Nat. Rev. Mater.* **3**, 5 (2018).
- [33] C. Suh, C. Fare, J. A. Warren, and E. O. Pyzer-Knapp, Evolving the materials genome: How machine learning is fueling the next generation of materials discovery, *Annu. Rev. Mater. Res.* **50**, 1 (2020).
- [34] E. O. Pyzer-Knapp, J. W. Pitera, P. W. J. Staar, S. Takeda, T. Laino, D. P. Sanders, J. Sexton, J. R. Smith, and A. Curioni, Accelerating materials discovery using artificial intelligence, high performance computing and robotics, *npj Comput. Mater.* **8**, 84 (2022).
- [35] A. Moradi Dakhel, V. Majdinasab, A. Nikanjam, F. Khomh, M. C. Desmarais, and Z. M. J. Jiang, GitHub Copilot AI pair programmer: Asset or liability? *J. Syst. Softw.* **203**, 111734 (2023).
- [36] S. Peng, E. Kalliamvakou, P. Cihon, and M. Demirer, The impact of AI on developer productivity: Evidence from GitHub Copilot, *arXiv:2302.06590*.
- [37] B. Yetiştiren, I. Özsoy, M. Ayerdem, and E. Tüzün, Evaluating the code quality of AI-assisted code generation tools: An empirical study on GitHub Copilot, Amazon CodeWhisperer, and ChatGPT, *arXiv:2304.10778*.
- [38] E. Mosqueira-Rey, E. Hernández-Pereira, D. Alonso-Ríos, J. Bobes-Bascarán, and A. Fernández-Leal, Human-in-the-loop machine learning: A state of the art, *Artif. Intell. Rev.* **56**, 3005 (2023).
- [39] F. Adams, A. McDannald, I. Takeuchi, and A. G. Kusne, Human-in-the-loop for Bayesian autonomous materials phase mapping, *Matter* **7**, 697 (2024).
- [40] Scientific autonomous reasoning agent (SARA): Integrating materials theory, experiment and computation—Research areas—AFOSR.
- [41] R. T. Bell, A. G. Jacobs, V. C. Sorg, B. Jung, M. O. Hill, B. E. Treml, and M. O. Thompson, Lateral temperature-gradient method for high-throughput characterization of material processing by millisecond laser annealing, *ACS Comb. Sci.* **18**, 548 (2016).
- [42] M.-C. Chang, S. Ament, M. Amsler, D. R. Sutherland, L. Zhou, J. M. Gregoire, C. P. Gomes, R. B. van Dover, and M. O. Thompson, Probabilistic phase labeling and lattice refinement for autonomous materials research, *npj Comput. Mater.* **11**, 148 (2025).
- [43] See Supplemental Material at <http://link.aps.org/supplemental/10.1103/c362-349b> for detailed information on the various performance metrics of SARA-H, experimental details, and phase diagrams of the materials systems, which includes Refs. [44–50].
- [44] M. C. Allison, M. Avdeev, S. Schmid, S. Liu, T. Söhnel, and C. D. Ling, Synthesis, structure and geometrically frustrated magnetism of the layered oxide-stannide compounds $\text{Fe}(\text{Fe}_{3-x}\text{Mn}_x)\text{Si}_2\text{Sn}_7\text{O}_{16}$, *Dalton Trans.* **45**, 9689 (2016).
- [45] S. Das and V. Jayaraman, SnO_2 : A comprehensive review on structures and gas sensors, *Prog. Mater. Sci.* **66**, 112 (2014).
- [46] J. A. Caraveo-Frescas, P. K. Nayak, H. A. Al-Jawhari, D. B. Granato, U. Schwingenschlögl, and H. N. Alshareef, Record mobility in transparent p-type tin monoxide films and devices by phase engineering, *ACS Nano* **7**, 5160 (2013).
- [47] Y. Masuda, H. Masumoto, A. Baba, T. Goto, and T. Hirai, Crystal growth, dielectric and polarization reversal properties of $\text{Bi}_4\text{Ti}_3\text{O}_{12}$ single crystal, *Jpn. J. Appl. Phys.* **31**, 3108 (1992).
- [48] A. Q. Jiang, Z. X. Hu, and L. D. Zhang, Investigations of morphotropic phase transformations in the solid solution of $\text{Bi}_4\text{Ti}_3\text{O}_{12}$ and $\text{Bi}_2\text{Ti}_4\text{O}_{11}$ accompanied by defect dipole orientation and oxygen vacancy migration, *J. Appl. Phys.* **85**, 1739 (1999).
- [49] M. Toyoda, Y. Hamaji, and K. T. Payne, Synthesis and characterization of $\text{Bi}_4\text{Ti}_3\text{O}_{12}$ thin films by sol-gel processing, *Jpn. J. Appl. Phys.* **32**, 4158 (1993).
- [50] T. Nakamura, R. Muhammet, M. Shimizu, and T. Shiosaki, Preparation of *c*-axis-oriented $\text{Bi}_4\text{Ti}_3\text{O}_{12}$ thin films by metalorganic chemical vapor deposition, *Jpn. J. Appl. Phys.* **32**, 4086 (1993).
- [51] N. Satoh, T. Nakashima, and K. Yamamoto, Metastability of anatase: Size dependent and irreversible anatase-rutile phase transition in atomic-level precise titania, *Sci. Rep.* **3**, 1959 (2013).
- [52] Z.-H. Cui, F. Wu, and H. Jiang, First-principles study of relative stability of rutile and anatase TiO_2 using the random phase approximation, *Phys. Chem. Chem. Phys.* **18**, 29914 (2016).
- [53] N. H. Vu, H. V. Le, T. M. Cao, V. V. Pham, H. M. Le, and D. Nguyen-Manh, Anatase-rutile phase transformation of titanium dioxide bulk material: A DFT + U approach, *J. Phys.: Condens. Matter* **24**, 405501 (2012).
- [54] W. Sun, S. T. Dacek, S. P. Ong, G. Hautier, A. Jain, W. D. Richards, A. C. Gamst, K. A. Persson, and G. Ceder, The thermodynamic scale of inorganic crystalline metastability, *Sci. Adv.* **2**, e1600225 (2016).
- [55] M. Amsler, V. I. Hegde, S. D. Jacobsen, and C. Wolverton, Exploring the high-pressure materials genome, *Phys. Rev. X* **8**, 041021 (2018).
- [56] D. A. H. Hanaor and C. C. Sorrell, Review of the anatase to rutile phase transformation, *J. Mater. Sci.* **46**, 855 (2011).
- [57] V. Borisenko and P. Hesketh, *Rapid Thermal Processing of Semiconductors* (Springer, Verlag, 1997).
- [58] S. Ament, M. Amsler, D. R. Sutherland, M.-C. Chang, D. Guevarra, A. B. Connolly, J. M. Gregoire, M. O. Thompson, C. P. Gomes, and R. B. van Dover, Autonomous materials synthesis via hierarchical active learning of nonequilibrium phase diagrams, *Sci. Adv.* **7**, eabg4930 (2021).
- [59] D. R. Sutherland, A. B. Connolly, M. Amsler, M.-C. Chang, K. R. Gann, V. Gupta, S. Ament, D. Guevarra, J. M. Gregoire, C. P. Gomes, R. Bruce van Dover, and M. O. Thompson, Optical identification of materials transformations in oxide thin films, *ACS Comb. Sci.* **22**, 887 (2020).
- [60] A. Kumar, V. Sindhwani, and P. Kambadur, Fast conical hull algorithms for near-separable non-negative matrix factorization, in *International Conference on Machine Learning* (PMLR, Atlanta, Georgia, 2013), pp. 231–239.
- [61] S. Ament, S. Daulton, D. Eriksson, M. Balandat, and E. Bakshy, Unexpected improvements to expected improvement for Bayesian optimization, in *Proceedings of the 37th Conference on Neural Information Processing Systems (NeurIPS)* (Neural Information Processing Systems Foundation, Inc., San Diego, CA, 2023), pp. 20577–20612.

- [62] D. Ackley, *A Connectionist Machine for Genetic Hillclimbing* (Springer Science & Business Media, New York, 2012), Vol. 28.
- [63] Q. Liang, A. E. Gongora, Z. Ren, A. Tiihonen, Z. Liu, S. Sun, J. R. Deneault, D. Bash, F. Mekki-Berrada, S. A. Khan, *et al.*, Benchmarking the performance of Bayesian optimization across multiple experimental materials science domains, *npj Comput. Mater.* **7**, 188 (2021).
- [64] H. A. Harwig, On the structure of bismuth sesquioxide: The α , β , γ , and δ -phase, *Z. Anorg. Allg. Chem.* **444**, 151 (1978).
- [65] N. Cornei, N. Tancrét, F. Abraham, and O. Mentré, New ϵ -Bi₂O₃ metastable polymorph, *Inorg. Chem.* **45**, 4886 (2006).
- [66] P. Shuk, H. D. Wiemhöfer, U. Guth, W. Göpel, and M. Greenblatt, Oxide ion conducting solid electrolytes based on Bi₂O₃, *Solid State Ion.* **89**, 179 (1996).
- [67] V. Fruth, G. Dobrescu, V. Bratan, C. Hornoiiu, S. Preda, C. Andronesu, and M. Popa, Structural and electrochemical features of Bi₂O₃-based fast oxide ion conductors, *J. Eur. Ceram. Soc.* **27**, 4421 (2007).
- [68] T. Takahashi, H. Iwahara, and T. Arao, High oxide ion conduction in sintered oxides of the system Bi₂O₃-Y₂O₃, *J. Appl. Electrochem.* **5**, 187 (1975).
- [69] I. Jeong, S. J. Jeong, B.-H. Yun, J.-W. Lee, C.-W. Lee, W. Jung, and K. T. Lee, Physically driven enhancement of the stability of Bi₂O₃-based ionic conductors via grain boundary engineering, *NPG Asia Mater.* **14**, 53 (2022).
- [70] R. T. Bell, P. A. Beaucage, M. J. Murphy, A. B. Connolly, U. Wiesner, D. Ginley, R. B. Van Dover, and M. O. Thompson, Rapid identification of synthetic routes to functional metastable phases using x-ray probed laser anneal mapping (XPLAM) time-temperature quench maps, *Chem. Mater.* **33**, 4328 (2021).
- [71] Y. F. Kargin, S. N. Ivicheva, and V. V. Volkov, Phase relations in the Bi₂O₃-TiO₂ system, *Russ. J. Inorg. Chem.* **60**, 619 (2015).
- [72] D. Sutherland, Autonomous methods enable discovery of metastable materials, Ph.D. thesis, Cornell University, 2023.
- [73] L. Liu, H. Zhao, J. M. Andino, and Y. Li, Photocatalytic CO₂ reduction with H₂O on TiO₂ nanocrystals: Comparison of anatase, rutile, and brookite polymorphs and exploration of surface chemistry, *ACS Catal.* **2**, 1817 (2012).
- [74] A. Di Paola, M. Bellardita, and L. Palmisano, Brookite, the least known TiO₂ photocatalyst, *Catalysts* **3**, 36 (2013).
- [75] H. Zhou, T.-J. Park, and S. S. Wong, Synthesis, characterization, and photocatalytic properties of pyrochlore Bi₂Ti₂O₇ nanotubes, *J. Mater. Res.* **21**, 2941 (2006).
- [76] R. E. Newnham, R. W. Wolfe, and J. F. Dorrian, Structural basis of ferroelectricity in the bismuth titanate family, *Mater. Res. Bull.* **6**, 1029 (1971).
- [77] H. S. Shulman, M. Testorf, D. Damjanovic, and N. Setter, Microstructure, electrical conductivity, and piezoelectric properties of bismuth titanate, *J. Am. Ceram. Soc.* **79**, 3124 (1996).
- [78] W. Wei, Y. Dai, and B. Huang, First-principles characterization of Bi-based photocatalysts: Bi₁₂TiO₂₀, Bi₂Ti₂O₇, and Bi₄Ti₃O₁₂, *J. Phys. Chem. C* **113**, 5658 (2009).
- [79] W. F. Yao, H. Wang, X. H. Xu, J. T. Zhou, X. N. Yang, Y. Zhang, and S. X. Shang, Photocatalytic property of bismuth titanate Bi₂Ti₂O₇, *Appl. Catal. A: Gen.* **259**, 29 (2004).
- [80] J. Hou, S. Jiao, H. Zhu, and R. V. Kumar, Bismuth titanate pyrochlore microspheres: Directed synthesis and their visible light photocatalytic activity, *J. Solid State Chem.* **184**, 154 (2011).
- [81] W.-F. Su and Y.-T. Lu, Synthesis, phase transformation and dielectric properties of sol-gel derived Bi₂Ti₂O₇ ceramics, *Mater. Chem. Phys.* **80**, 632 (2003).
- [82] C. L. Gomez, O. Depablos-Rivera, J. C. Medina, P. Silva-Bermudez, S. Muhl, A. Zeinert, and S. E. Rodil, Stabilization of the delta-phase in Bi₂O₃ thin films, *Solid State Ionics* **255**, 147 (2014).
- [83] S. Durmuş, V. Çorumlu, T. Çifci, I. Ermis, and M. Arı, Electrical, structural and thermal properties of nanoceramic (Bi₂O₃)_{1-x-y}(Ho₂O₃)_x(Tm₂O₃)_y ternary system, *Ceram. Int.* **39**, 5241 (2013).
- [84] C. D. Ling, R. L. Withers, S. Schmid, and J. G. Thompson, A review of bismuth-rich binary oxides in the systems Bi₂O₃-Nb₂O₅, Bi₂O₃-Ta₂O₅, Bi₂O₃-MoO₃, and Bi₂O₃-WO₃, *J. Solid State Chem.* **137**, 42 (1998).
- [85] S.-H. Jung, E. D. Wachsman, and N. Jiang, Structural stability and conductivity of cubic (WO₃)_x - (Dy₂O₃)_y - (Bi₂O₃)_{1-x-y}, *Ionics* **8**, 210 (2002).
- [86] M. Struzik, X. Liu, I. Abrahams, F. Krok, M. Malys, and J. R. Dygas, Defect structure and electrical conductivity in the pseudo-binary system Bi₃TaO₇-Bi₃NbO₇, *Solid State Ionics* **218**, 25 (2012).
- [87] E. I. Speranskaya, I. S. Rez, and L. V. Kozlova, The system bismuth oxide-titanium dioxide, *Izv. Akad. Nauk SSSR, Neorg. Mater.* **1**, 232 (1965).
- [88] A. A. Maier, V. A. Lomonov, V. A. Balashov, and N. G. Gorashchenko, Physicochemical foundations of the technology of single crystals with sillenite and eulytine structures, in *Research in the Field of Chemistry and Chemical Technology of Materials for the Electronics Industry* (1981), pp. 16–27.
- [89] L. Yao, S. Samantray, A. Ghosh, K. Roccapriore, L. Kovarik, S. Allec, and M. Ziatdinov, Operationalizing serendipity: Multi-agent AI workflows for enhanced materials characterization with theory-in-the-loop, *arXiv:2508.06569*.
- [90] Y. Du, B. Yu, T. Liu, T. Shen, J. Chen, J. G. Rittig, K. Sun, Y. Zhang, A. Krishnan, Y. Zhang, *et al.*, Accelerating scientific discovery with autonomous goal-evolving agents, *arXiv:2512.21782*.
- [91] Physical-chemical vapor deposition systems & coated products, <https://angstromengineering.com>.
- [92] M.-C. Chang, Enabling autonomous high-throughput XRD-based phase discovery in thin film systems, Ph.D. thesis, Cornell University, 2025.
- [93] M. Abel, T. Beecham, S. Graham, S. P. Kearney, J. R. Serrano, and L. M. Phinney, Noncontact surface thermometry for microsystems: LDRD final report, Technical Report SAND2006-6369, Sandia National Laboratories (SNL), Albuquerque, NM, and Livermore, CA (United States), 2006.
- [94] L. Smieska, K. A. Page, B. Ree, B. Zheng, H. Koerner, and A. R. Woll, The functional materials beamline at CHESS, *Synchrotron Radiat. News* **36**, 4 (2023).
- [95] G. Ashiotis, A. Deschildre, Z. Nawaz, J. P. Wright, D. Karkoulis, F. E. Picca, and J. Kieffer, The fast azimuthal in-

- tegration Python library: pyFAI, *J. Appl. Crystallogr.* **48**, 510 (2015).
- [96] C. Guo, G. Pleiss, Y. Sun, and K. Q. Weinberger, On calibration of modern neural networks, in *International Conference on Machine Learning* (PMLR, Sydney, Australia, 2017), pp. 1321–1330.
- [97] T. M. Hamill, Reliability diagrams for multicategory probabilistic forecasts, *Weather Forecast.* **12**, 736 (1997).
- [98] C. K. Williams and C. E. Rasmussen, *Gaussian Processes for Machine Learning* (MIT Press, Cambridge, MA, 2006), Vol. 2.
- [99] A. McHutchon and C. Rasmussen, Gaussian process training with input noise, in *Proceedings of the Advances in Neural Information Processing Systems*, edited by J. Shawe-Taylor, R. Zemel, P. Bartlett, F. Pereira, and K. Q. Weinberger (Curran Associates, Inc., Red Hook, NY, 2011), Vol. 24, pp. 1341–1349.
- [100] J. Tellinghuisen, Statistical error propagation, *J. Phys. Chem. A* **105**, 3917 (2001).
- [101] S. Ament, A. Witte, N. Garg, and J. Kusuma, Sustainable concrete via Bayesian optimization, [arXiv:2310.18288](https://arxiv.org/abs/2310.18288).
- [102] D. C. Liu and J. Nocedal, On the limited memory BFGS method for large scale optimization, *Math. Program.* **45**, 503 (1989).
- [103] W. H. Press, S. A. Teukolsky, W. T. Vetterling, and B. P. Flannery, *Numerical Recipes: The Art of Scientific Computing*, 3rd ed. (Cambridge University Press, Cambridge, 2007).
- [104] M.-C. Chang, M. Amsler, D. Sutherland, S. Ament, K. Gann, L. Zhou, L. Smieska, A. Woll, J. Gregoire, C. Gomes, R. van Dover, and M. Thompson, Data from: Autonomous materials exploration by integrating automated phase identification and AI-assisted human reasoning. [dataset] Cornell University Library eCommons Repository, 2026, <https://doi.org/10.7298/xrck-2516>.