MATERIALS SCIENCE

Autonomous materials synthesis via hierarchical active learning of nonequilibrium phase diagrams

Sebastian Ament[1]†, Maximilian Amsler[2,3]*†, Duncan R. Sutherland[2], Ming-Chiang Chang[2], Dan Guevarra[4], Aine B. Connolly[2], John M. Gregoire[4], Michael O. Thompson[2], Carla P. Gomes[1]*, R. Bruce van Dover[2]

Autonomous experimentation enabled by artificial intelligence offers a new paradigm for accelerating scientific discovery. Nonequilibrium materials synthesis is emblematic of complex, resource-intensive experimentation whose acceleration would be a watershed for materials discovery. We demonstrate accelerated exploration of metastable materials through hierarchical autonomous experimentation governed by the Scientific Autonomous Reasoning Agent (SARA). SARA integrates robotic materials synthesis using lateral gradient laser spike annealing and optical characterization along with a hierarchy of AI methods to map out processing phase diagrams. Efficient exploration of the multidimensional parameter space is achieved with nested active learning cycles built upon advanced machine learning models that incorporate the underlying physics of the experiments and end-to-end uncertainty quantification. We demonstrate SARA's performance by autonomously mapping synthesis phase boundaries for the Bi_2O_3 system, leading to orders-of-magnitude acceleration in the establishment of a synthesis phase diagram that includes conditions for stabilizing δ-Bi_2O_3 at room temperature, a critical development for electrochemical technologies.

INTRODUCTION

Artificial intelligence (AI) holds great promise for revolutionizing scientific fields as varied as biology (1), chemistry (2), physics (3), and economics (4). Much of AI's impressive recent successes have been in data-rich applications: AlphaFold (1), for example, uses a library of tens of thousands of existing protein data to build a highly successful model for protein folding. In fields that lack comparably vast data, however, a great opportunity lies in guiding the exploratory process itself to minimize the number of experiments that are required to achieve insights, i.e., active learning (AL) (5), and thereby accelerate the pace of scientific discovery (6).

These AI-guided efforts have shown great promise in the design of quantum experiments (7), drug development (8), and wind turbine control (9) and are of particular importance in materials research aimed at designing and optimizing functional materials that lie at the core of technological advances. High-throughput (HT) experimental synthesis and characterization of materials systems through thin-film deposition of inorganic composition spreads (10), so-called libraries, present a promising avenue to rapidly explore a vast chemical, structural, and property space (11). These methods have been well established for comprehensive synthesis of composition spaces with two to four components, where the resulting tens to thousands of materials can be evaluated via automated characterization. While this approach has been quite effective for identification of materials with desired properties, the opportunity for broader materials exploration to enable new technologies is highlighted by the limited exploration of synthesis conditions to date. The portion of the materials search space that has been explored is vanishingly small when considering the dynamic range of thermal processing conditions, which are inherent to processing-composition-structure-property (PCSP) relationships. The breadth of relevant thermal processing conditions makes exhaustive sampling untenable, and as HT experiments lower the time for an experiment cycle, human intervention in decision-making becomes a worsening bottleneck. Therefore, AI and AL are critical in both reducing the number of experiments to a more tractable scale and accelerating the process of making decisions to match the rate of incoming data (12–14).

Recently, HT experimentation and AL techniques have been combined in a closed-loop fashion, where an AI instance iteratively proposes a sequence of experiments to explore and discover new materials. These efforts include identifying phase-change materials via Bayesian AL (15), the discovery of NiTi-based shape memory alloys with low thermal hysteresis (16), the synthesis of $BaTiO_3$-based piezoelectrics with large electrostrain (17), the selective growth of carbon nanotubes (18,19), the search for perovskite-type materials for photovoltaic applications (20) and inorganic quantum dots (21), maximizing hole mobility of organic solar cells (22), and accelerating toughness optimization in additive manufacturing (23).

Despite this progress, the current state of the art exhibits considerable limitations. Chiefly, most attempts at closed-loop cycles still rely heavily on human intervention, preventing them from reaching true autonomy in materials discovery. Furthermore, although AL guidance has recently been deployed to great effect for discovery of optical phase-change thin-film materials (15), the search space was limited to presynthesized compositions using a single processing condition. The time and temperature scales relevant to nonequilibrium thermal processing of solid-state inorganic materials pose substantial problems for incorporating synthesis in an autonomous loop, although the utility of spanning synthesis and characterization in an AL framework has been demonstrated by several recent autonomous workflows for chemical synthesis (22, 24–26). The complexity and

[1]Department of Computer Science, Cornell University, Ithaca, NY 14853, USA. [2]Department of Materials Science and Engineering, Cornell University, Ithaca, NY 14853, USA. [3]Department of Chemistry and Biochemistry, University of Bern, Freiestrasse 3, CH-3012 Bern, Switzerland. [4]Division of Engineering and Applied Science, California Institute of Technology, Pasadena, CA 91125, USA.
*Corresponding author. Email: amsler.max@gmail.com (M.A.); gomes@cs.cornell.edu (C.P.G.)
†These authors contributed equally to this work.

the degrees of freedom of the PCSP space are particularly challenging to incorporate in autonomous experimentation when considering metastable materials that form far from equilibrium at different, often unpredictable processing conditions (27,28). Commonly deployed off-the-shelf AL models are often not sufficient for achieving highly efficient learning and are frequently outperformed by random search with twice the number of samples (29), a problem that is exacerbated with increased dimensionality of the search space (30).

Expert human scientists navigate complex search spaces by incorporating their prior knowledge, such as physics-based models that underlie the acquired data. Incorporating such knowledge in AL often requires the development of new AI methods. Lastly, exploration via AL critically relies on uncertainty quantification in the not-yet-sampled regions of parameter space, which, for complex experimental workflows, requires error propagation. Arguably, the most immediate obstacle to accelerating experimental exploration via AL lies in the dual challenges of developing noise models for each type of experiment and integrating them into a computational framework for end-to-end uncertainty quantification. In aggregate, these challenges motivate the establishment of a framework that integrates AI methods at multiple scales to perform scientifically meaningful interpretation, modeling, and uncertainty quantification of multiple streams of incoming data.

Our vision of the Scientific Autonomous Reasoning Agent (SARA) (31) is to develop a fully autonomous HT materials discovery and exploration framework by integrating robotic HT materials synthesis (26) with AI instances to accelerate both materials synthesis and analysis. In particular, SARA aims to automate the representation, characterization, planning, optimization, and learning of materials knowledge in a fully integrated manner. To achieve this goal, we envision the deployment of agents, which individually specialize on specific subtasks but closely interact with each other to accelerate the discovery efforts. These agents include, but are not limited to, synthesis and probing robotics to conduct experiments and highly optimized, physics-based AI models that evaluate currently available data with their associated uncertainties and that drive AI-guided discovery.

In this work, we take strides toward realizing this vision and present a fully integrated, autonomous framework that iteratively maps out the synthesis phase boundaries of metastable compounds in a closed-loop fashion. To this end, we incorporate a system of nested (32) cycles harnessed by SARA's specialized AI agents to synthesize and explore thin-film libraries with lateral gradient laser spike annealing (lg-LSA) (33): An internal (highest frequency) autonomous loop iteratively proposes optimized property measurements of a given lg-LSA stripe using a hierarchy of optical characterization techniques, while an external autonomous loop proposes and executes the next lg-LSA synthesis via a model that aggregates knowledge obtained by inner loop iterations. This architecture can be readily expanded and nested into higher-level loops that, e.g., optimize thin-film deposition, materials systems, and quantum materials computation.

SARA's nested synthesis, microscopy imaging, and reflectance spectroscopy loops driven by the specialized AIs with AL reflect the hierarchical nature of scientific discovery. A primary goal of studying PCSP relationships is the enumeration of all possible syntheses that yield unique materials, a knowledge base that must be built from synthesis phase diagrams over a broad range of synthesis techniques, multiple parameter spaces defined within each technique, and

many experimental campaigns to map synthesis phase diagrams in those spaces. Coordination among the levels of hierarchy is critical for maximizing high-level knowledge generation from low-level experiments, which guides our development of nested AL algorithms that seamlessly incorporate task coordination and uncertainty propagation. This framework is extensible with respect to incorporation of additional levels of hierarchy and/or expansion of techniques, such as additional property measurements and on-the-fly quantum mechanical calculations (34, 35), that enrich the knowledge within a given level of hierarchy. Networking of capabilities and knowledge sources elevates the use of AI and AL from process optimization to accelerated scientific discovery, a grand vision of AI-assisted science.

RESULTS

Our goal is to explore synthesis phase diagrams, especially the relatively unexplored ultrafast-annealing region where metastable polymorphs of metal oxides are more likely to form. These metastable oxide materials often exhibit improved properties over thermodynamic ambient ground states and are relevant for countless industrial applications. The cubic high-temperature polymorph of ZrO_2, for example, is frequently used as a thermal coating material (36–38) because of its low thermal conductivity, while the anatase phase of TiO_2 has attracted interest as a photocatalytic material (39–41). These are only two of the most prominent examples of materials systems where metastable phases outperform their respective equilibrium counterparts.

Here, we study the Bi-O system, which exhibits a rich phase diagram with dozens of experimentally observed polymorphs. In particular, we focus on the Bi_2O_3 composition, for which five distinct crystalline phases are known (42, 43). The monoclinic α-Bi_2O_3 is the thermodynamic ground state at room temperature, while four high-temperature phases have been reported: tetragonal β-Bi_2O_3, body-centered cubic γ-Bi_2O_3, cubic δ-Bi_2O_3, and orthorhombic ϵ-Bi_2O_3. The metastable δ phase has attracted interest as a solid oxide electrolyte in fuel cells (44): Because of its defective fluorite-type crystal structure with a high concentration of oxygen vacancies, δ-Bi_2O_3 has the highest oxygen ion conductivity of any solid oxide known to date. Unfortunately, it exhibits only a narrow thermodynamic stability window between 727° and 824°C, which has so far precluded its use on an industrial scale. Substitution of yttrium or rare earth oxides can stabilize δ-Bi_2O_3 to room temperature but leads to a degraded ion conductivity. Hence, efforts have been aimed at finding routes to retain phase-pure δ-Bi_2O_3 to ambient conditions (45).

Our samples are deposited as amorphous thin films by reactive sputtering on a silicon substrate. For other materials systems, composition spreads can be similarly deposited, allowing the mapping of a composition gradient $c(x)$ to the location x on the substrate. We process the thin-film libraries using lg-LSA to form and kinetically trap metastable phases during the quench to ambient conditions. In contrast to conventional methods for annealing thin-film samples, such as hot plate, furnace, and rapid thermal annealing (46), lg-LSA allows a controlled and rapid thermal processing over a wide range of conditions in a spatially confined region of less than 1 mm, with quench rates of 10^4 to 10^7 K/s and peak temperatures T_p up to 1400°C (limited by melt of the silicon substrates). Scanning a laser beam with a bi-Gaussian–like power profile (see the backdrop in the left panel of Fig. 1) over the film allows a single lg-LSA stripe to
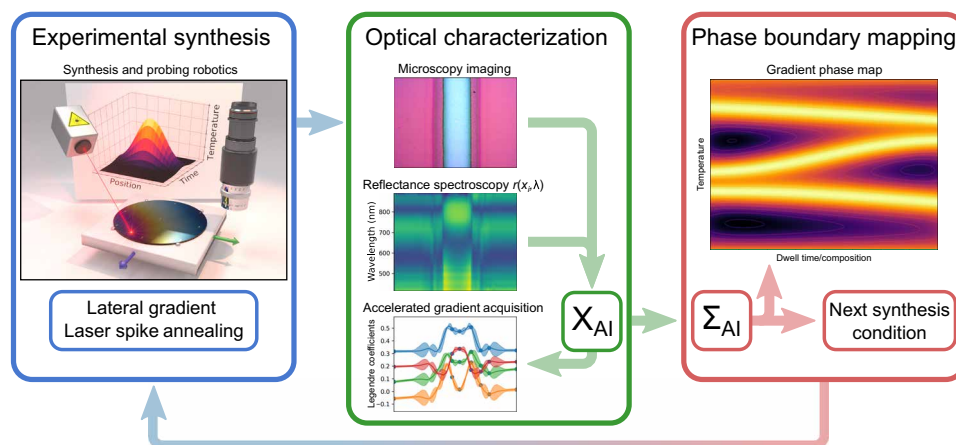
Fig. 1. SARA's closed-loop autonomous materials synthesis and discovery cycle. Starting from a set of initially selected processing conditions, SARA synthesizes an lg-LSA stripe on a thin-film library and subsequently sends it to its characterization AI agent, X_{AI}. (**Left**) Schematic illustration of the lg-LSA/camera setup with the bi-Gaussian power profile in the backdrop, laser to the left, camera to the right, and thin-film sample mounted on a stage. Using a hierarchy of characterization techniques, X_{AI} analyzes the stripe to determine intricate changes in its optical properties. In particular, X_{AI} first acquires a microscope image to determine the positions of likely phase boundaries, which informs the reflectance spectroscopy measurements. X_{AI}'s physics-informed AL model accelerates the spectroscopy acquisition, resulting in an accurate gradient model of the lg-LSA stripe. (**Middle**) Microscope image, reflectance spectroscopy heatmap, and first four Legendre coefficients from the X_{AI} representation for a representative lg-LSA stripe. Lastly, the gradients are fed into SARA's synthesis AI agent, Σ_{AI}, which generates a gradient phase boundary map and also proposes the next experimental processing conditions to improve the phase boundary with as few experiments as possible. (**Right**) Model gradient phase map showing high-gradient regions in yellow.

produce a spatially inhomogeneous thermal profile $T_{T_p, \tau}(x)$ (where x runs across the stripe). The duration of heating is characterized by a dwell time τ defined by the ratio of the laser full width at half maximum (FWHM) divided by the scan velocity of the laser (typical dwells range from 100 to 10,000 μs). Hence, at a given dwell time τ, a single lg-LSA experiment produces a continuous range of temperature conditions wherein phase transitions, including formation of the sought metastable phases, need to be detected with a speed and level of automation commensurate with this robotic synthesis procedure to fully capitalize and elevate high-throughput synthesis to high-throughput discovery of phase boundaries.

To reduce both computational and experimental cost, we need to autonomously map out the processing phase space $\{x, \tau, T_p\}$ with as few synthesis experiments and property measurements as possible. Because the lg-LSA is an irreversible method, a specific position x [and potentially its associated composition $c(x)$ in the presence of a composition gradient] can only be annealed once, further emphasizing the need for optimizing the selection of the processing conditions. Once an lg-LSA stripe is processed, a conclusive structural characterization across the thermal gradient is possible with grazing-incidence high-intensity x-ray diffraction (XRD) to resolve the crystal structure (*47*). However, access to synchrotron facilities capable of producing x-rays with appropriate wavelength, intensity, and micrometer-scale spatial resolution comprises an inherently limited resource that motivates development of alternative phase boundary detection methods. To address this issue, we developed a complementary technique based on microscopy imaging and optical spectroscopy to rapidly assess phase boundaries. We recently demonstrated that structural phase changes are directly associated with changes in the optical thin-film properties of transparent films, in particular the optical thickness nd (*47*), where n is the refractive index and d is the film thickness. Essentially, the gradients of the optical measurements across an lg-LSA stripe provide a means to

map out phase boundaries without explicit crystallographic phase identification, thereby producing an unlabeled processing phase diagram without costly XRD experiments.

Here, we put forth how SARA integrates lg-LSA synthesis and optical phase boundary detection in a hierarchical autonomous workflow by using characterization and synthesis agents, X_{AI} (pronounced chi AI) and Σ_{AI} (pronounced sigma AI), respectively, as illustrated in Fig. 1. Starting with an initial processing condition, SARA synthesizes an lg-LSA stripe on a thin-film library. Then, SARA uses its internal characterization agent X_{AI} to probe the stripe using a set of optical techniques: (i) microscopy imaging to rapidly inspect the anneal stripe (see the top panel in "Optical characterization" in Fig. 1) and (ii) more elaborate, but costly, reflectance measurements (see "Reflectance spectroscopy" in Fig. 1). In particular, X_{AI} uses the observed features from the micrograph as prior knowledge to guide and acquire an accurate reflectance map with as few measurements as possible. The gradients of the reflectance map are then fed into SARA's synthesis AI agent Σ_{AI}, which incorporates the reflectance gradient information of each lg-LSA stripe into a phase boundary map as a function of the parameters $\{x, \tau, T\}$. The high-gradient regions of this map determine the boundaries between phase fields and produce an unlabeled processing phase diagram (see "Phase boundary mapping" in Fig. 1). Σ_{AI} is also responsible for proposing the next most promising synthesis conditions to effectively explore the search space. We discuss X_{AI} and Σ_{AI} in detail below.

X_{AI}: Accelerating data acquisition and characterization
X_{AI}'s primary task is to construct an accurate reflectance spectroscopy map $r(x, \lambda)$ of an lg-LSA sample annealed at T_p and τ while measuring it at as few positions x_i across the stripe as possible. Because the acquisition time for a single such measurement $r(x_i, \cdot)$ is around 4.5 s, an exhaustive scan across a stripe of 1.5 mm in

10–µm intervals requires more than 11 min, forming one of the main bottlenecks of our HT experimental setup. To accelerate the reflectance data acquisition, we propose an AL scheme that takes advantage of multimodal measurements and incorporates physical structure into a Gaussian process (GP) regression model to yield highly optimized data acquisition and analysis.

The overall workflow of the X_{AI} cycle is outlined in Fig. 2A. In a first step, SARA captures a microscope image of an lg-LSA stripe to analyze the overall condition of the anneal and to extract key features. This single RGB (red-green-blue) image of a stripe is inherently throughput-matched to the lg-LSA synthesis, producing prior knowledge for the X_{AI}'s AL cycle to accelerate reflectance measurements. A representative microscope image is shown in Fig. 2B. These micrographs can be used to rapidly assess the conditions and the integrity of the anneal. Obvious damage of the thin film such as delamination and scratches or contamination such as dust particles, residual lithography artifacts, and dirt can be easily detected, which invalidates the lg-LSA stripe and can trigger resynthesis. The incorporation of such automated quality control in the autonomous loop alleviates responsibility for the X_{AI} loop to effectively respond to invalid data, a critical aspect of autonomous workflows for robust operation (13).

SARA proceeds by constructing a stripe-specific GP kernel that incorporates the underlying physics of both the lg-LSA and optical spectroscopy processes. Notably, the bi-Gaussian power profile produces stripes of nearly perfect lateral symmetry at steady state, with their centers reaching the corresponding peak temperatures T_p and the continuous variation in lateral thermal gradient mirrored on each side of the stripe. We incorporate this structure into the kernel of X_{AI} by forcing its main component to be symmetric around the center of a stripe (see the "X_{AI}" section in the "materials and methods" part). In addition, SARA extracts key features of the stripe

texture from the micrograph to further improve the kernel design, i.e., by identifying the stripe center and by detecting systematic optical changes across the stripe that we associate with structural transitions (47, 48). These optical transitions are identified by peaks in the gradient signal across a stripe, the locations of which are shown as vertical yellow lines in Fig. 2B. Furthermore, the two outermost detected peaks in the gradient signal give an estimate of how wide the lg-LSA stripe is, i.e., where the unannealed, amorphous film ends and the crystallization begins. We use slightly broadened peaks in the RGB gradient signal (purple line in Fig. 2C) and the overall width of the lg-LSA stripe (red line in Fig. 2C) as the RGB and LSA prior, respectively. These two functions are then used to rescale the kernel of the GP in the X_{AI} cycle. Lastly, we account for small thickness variations of the film across the stripe by adding a linear component to the kernel.

To improve the efficiency of X_{AI}, the reflectance $r(x, \lambda)$ at any position x is expanded in Legendre polynomials as a function of wavelength λ before it is fed into the GP. Because the reflectance varies smoothly with λ, the Legendre expansion can be truncated between the 10th and 20th order at essentially no loss in accuracy (49) (see fig. S1), which reduces the dimensionality from the 2046 measured photon wavelength to a compact space of 10 to 20 Legendre coefficients. For our system, we use 16 coefficients throughout. Figure 1 (bottom middle) shows the first four Legendre coefficients of the reflectance data and our GP model's posterior predictive mean and uncertainty for those coefficients.

To demonstrate the advantage of our specialized X_{AI} kernel with respect to a set of conventional kernels, we perform statistical benchmarks on 617 lg-LSA experiments at distinct conditions, $(T_p, \tau)_j$. For each of the stripes, we measure the reflectance at n randomly selected positions $\{x_i\}_j^n$ on a grid spaced 10 µm apart and use these measurements as inputs to a GP model with different kernels. The
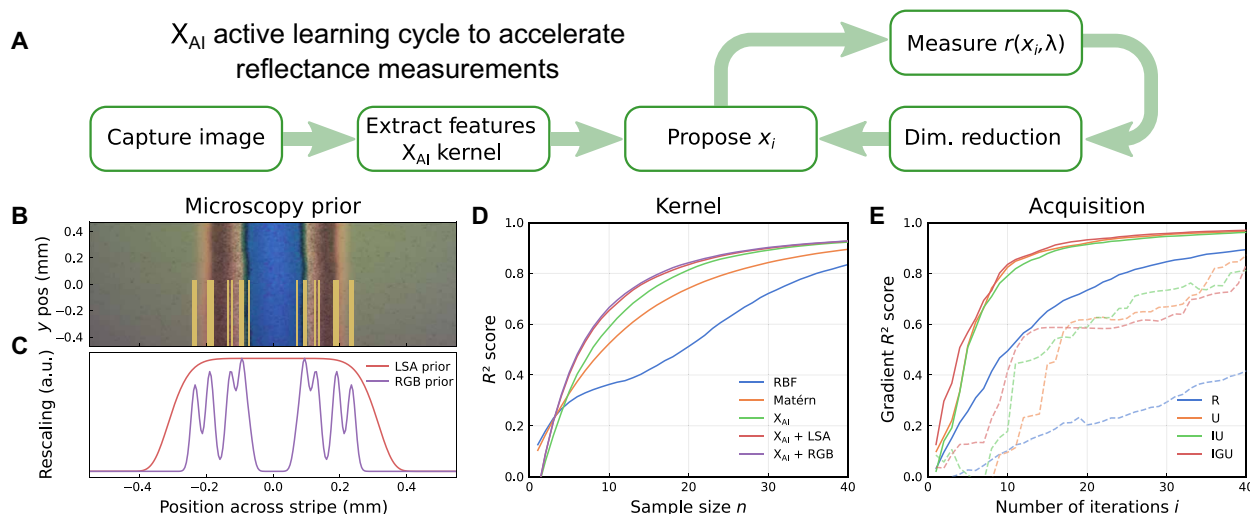

Fig. 2. The characterization AL loop to accelerate acquisition of the reflectance spectroscopy necessary for phase boundary detection. (**A**) Overall workflow. A microscope image (**B**) is captured to extract the stripe features, which are fed as a scaling function to the X_{AI} kernel. The core features are the LSA prior and the RGB transition prior, which are sums of generalized Gaussian functions, as shown in (**C**). The corresponding gradient peak positions are denoted as yellow vertical lines in (B). X_{AI} takes these functions as prior knowledge to set up a stripe-specific kernel that facilitates rapid model convergence. The AL loop is performed iteratively on reflectance measurements $r(x_i, \lambda)$ over positions x_i, which are expanded into Legendre polynomials to reduce the dimensionality (see the middle panel in Fig. 1). (**D**) Performance of different kernel designs, illustrating that our X_{AI} kernel with both LSA (X_{AI} + LSA) and LSA + RGB (X_{AI} + RGB) priors outperforms other conventional kernels. (**E**) Performances of the different acquisition function (R, random; U, uncertainty; IU, integrated uncertainty; IGU, integrated gradient uncertainty sampling). The solid lines represent the X_{AI} + RGB kernel, while the dashed lines correspond to the RBF kernel.

ground truth is exhaustively measured across the whole stripe ranging over 1.5 mm, corresponding to a total of 151 measurements. For a range of n, we repeat this test 32 times for every stripe with independent random locations and average the coefficient of determination R^2 for each kernel on the exhaustive data. This reduces the statistical noise in the results to a negligible value. Furthermore, we benchmark every kernel with a range of length scales and select the best in terms of R^2 score (see the "Error metrics" section). By construction, our benchmark disentangles the effects of AL and kernel design, and the kernel with the right inductive bias will express the data best, even if all measurements are random. The results of this benchmark are illustrated in Fig. 2D, showing the performance of the various kernels with respect to the number of random measurements n. The radial basis function (RBF) kernel performs poorly, barely reaching an R^2 score of 0.8 within 37 measurements. The Matérn kernel performs better, requiring $n = 25$ to reach the same score. The X_{AI} kernels perform best: Depending on whether prior knowledge from the microscope image is included ("X_{AI} + LSA" with LSA prior only and "X_{AI} + RGB" with LSA and gradient peak prior), we obtain an R^2 value of 0.8 with as few as 16 sampling points. The precise modeling of the optical measurements, its incorporation into the AL model, and the model's initialization with the RGB image prior knowledge all contribute to the fast learning rate at the onset of autonomous experimentation, as required for efficient AL.

Having designed the kernel for the X_{AI} cycle, we turn our attention to the acquisition function, that is, the function that chooses the next measurement based on the available information. An important component of many performant acquisition functions is the reduction of uncertainty in a target variable. Here, we benchmark three different acquisition functions, two of which are nonstandard. In particular, we study uncertainty (U) sampling, which chooses the next measurement at the point of maximum uncertainty in the Legendre coefficients; integrated uncertainty (IU) sampling, which selects the point that minimizes the integrated uncertainty over the whole sampling domain; and integrated gradient uncertainty (IGU) sampling, which is similar to IU but reduces the overall uncertainty in the gradients of the model. The last strategy targets our quantity of interest, because the reflectance gradients are indicative of the phase boundaries in the processing phase diagram. For this reason, we quantify the error of the model in the gradients, rather than the error to the observed data. Because we cannot directly observe the gradients, we generate ground truth data by training our GP model on the exhaustive measurements and taking the derivative of the fitted model. We then record the R^2 score of the derivatives of the model for each of the acquisition functions at every iteration.

In Fig. 2E, we show the performance of the various sampling strategies as a function of AL iteration i, using either the X_{AI} + RGB kernel (solid lines) or the RBF kernel (dashed lines). The best performance is achieved with the stripe-specific, highly optimized X_{AI} + RGB kernel in conjunction with IGU sampling, reaching an R^2 score of 0.8 and 0.9 within 9 and 15 iterations, respectively. Note that random sampling with the best kernel design still outperforms the best sampling strategies with the worst kernel. Furthermore, the acquisition functions do not differ markedly with the X_{AI} + RGB kernel, highlighting the importance of incorporating the problem structure into our AI model and AL cycle. Compared to random sampling with an RBF kernel, the best strategy accelerates the acquisition and characterization by a factor of 9.7 for an R^2 of 0.8, approximately one order of magnitude.

Σ_{AI}: Accelerating phase exploration and processing conditions

Once an lg-LSA stripe has been processed by X_{AI}, its output reflectance gradient information is fed into the external synthesis AI agent, Σ_{AI}. Its main task is threefold: (i) assemble the incoming data, (ii) propagate uncertainty from every lg-LSA experiment to predict the gradient signal and its uncertainty throughout the search space, and (iii) ultimately propose new conditions for the synthesis experiments. The overall workflow of this process, which integrates the techniques described below, is shown in Fig. 3A.

The optical data of an lg-LSA anneal are processed through the nested X_{AI} loop, the output of which is the gradients of the reflectance spectroscopy across a stripe, $g(x) = \|\partial_x r(x,\cdot)\|_2$. This spatial gradient information is then transformed onto a temperature scale based on the Gaussian-type temperature profile $T_{T_p,\tau}(x)$ shown in Fig. 3B (blue line). Because the X_{AI} kernel is symmetric up to the linear term, the gradient information is symmetric about the peak temperature T_p so that we only need to sample $g(x)$ along one side from the stripe center (orange crosses in Fig. 3B).

In principle, one single lg-LSA stripe would produce the complete temperature conditions between room temperature T_r and T_p at a given dwell time τ. Hence, the set of metastable materials and their transition conditions would be available from a single stripe if one selected a high T_p (e.g., 1400°C) and $T_{min} = T_r$. In practice, the concomitant increase in temperature gradient with T_p would require progressively higher spatial resolution to characterize the full range of transitions and results in undesirably high uncertainty in the modeled temperature. With our experimental characterization technique, the spatial resolution is limited to approximately 10 μm, and thus, the design of lg-LSA synthesis conditions must be done under consideration of the position-dependent temperature variation within a single spectroscopy measurement, which makes the selection of T_p at a given τ a nontrivial decision based on the aggregate information that can be gained from the entire lg-LSA stripe.

Properly propagating the multiple sources of uncertainty from synthesis and characterization through the model of the phase boundary map is extremely important: In standard GP regression, the inputs are assumed to be free of noise, but accounting for such errors is crucial when dealing with experimental measurements. Here, we include and propagate the uncertainties of the inputs due to two sources. First, the peak temperature reached in an lg-LSA anneal can vary within an error range of up to $\sigma_{T_p} = 25°C$ at 1400°C due to fluctuation in the laser power, even after reaching steady state (Peak error in Fig. 3B). Second, the temperature profile itself gives rise to an error proportional to the spatial rate of change $\sigma_T(x) \propto |\partial_x T_{T_p,\tau}(x)|$, as shown by the green area in Fig. 3B. Note that the error bars in Fig. 3B are not to scale and intended solely for a schematic illustration.

As opposed to the X_{AI} model of optical spectroscopy, the gradient map in synthesis space has no analogous physics-based model, in part because too few synthesis phase diagrams are known, and their underlying features remain an open question. The large dynamic range of dwell time motivates its logarithmic sampling, and the distinct influence of temperature and dwell time on synthesis motivates independent parameterization of these two dimensions. While we aim to learn more structured representations of synthesis phase diagrams in future refinements of the SARA framework, for the purposes of the present work, we find a Matérn kernel, with separate length scales for the temperature and dwell time dimensions, to
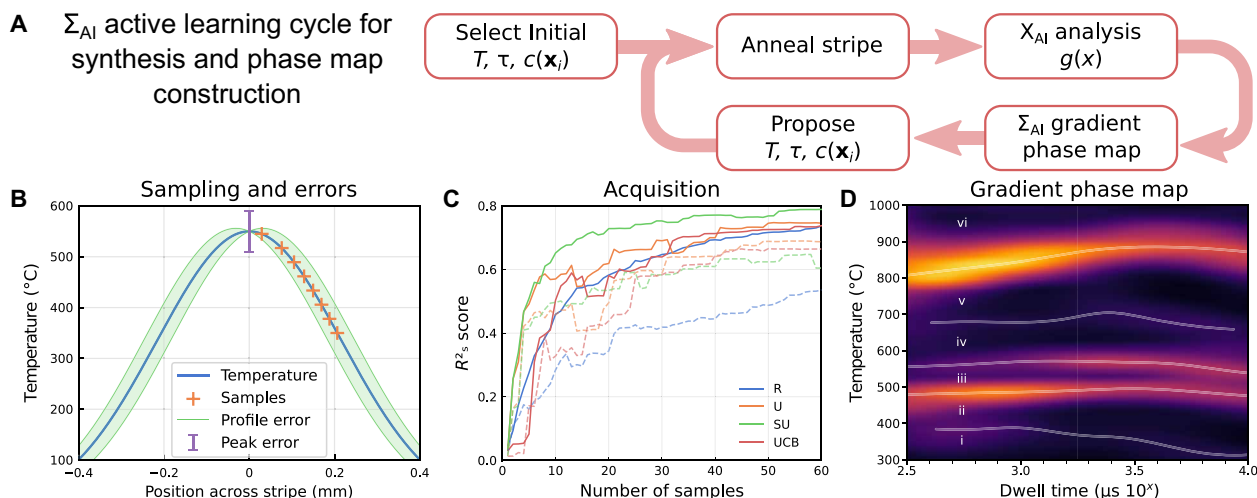
A Σ_{AI} active learning cycle for synthesis and phase map construction

Select Initial T, τ, $c(\mathbf{x}_i)$ → Anneal stripe → X_{AI} analysis $g(x)$ → Σ_{AI} gradient phase map → Propose T, τ, $c(\mathbf{x}_i)$ →

Fig. 3. The synthesis AL loop to accelerate materials exploration. (**A**) Overall external workflow. Starting from an initial set of conditions, an lg-LSA stripe is annealed and processed by X_{AI}. The gradients are then fed into the Σ_{AI} agent, which constructs a (preliminary) gradient phase map and proposes the next experimental conditions. (**B**) The transformation of the X_{AI} reflectance gradients requires a rigorous error assessment and propagation. Because of the symmetric X_{AI} kernel, only one side from the stripe center is sampled on a uniform temperature mesh. The errors propagated to the Σ_{AI} stem from the variation in the peak temperature (peak error) and the gradient of the temperature profile (profile error). (**C**) Performance of different Σ_{AI} acquisition strategies: random (R), uncertainty (U), stripe uncertainty (SU), and upper confidence bound (UCB) sampling. The solid and dashed lines correspond to GP regression with and without input uncertainty, respectively. (**D**) Gradient phase map of Bi_2O_3, where the peak ridges are highlighted with light lines. The phase regions are labeled a posteriori with selected XRD measurements, from low to high temperatures: (i) amorphous as-deposited; (ii) rearranged, densified amorphous; (iii) δ-Bi_2O_3; (iv) mixed-phase region of δ-Bi_2O_3 and β-Bi_2O_3; (v) pure β-Bi_2O_3; and, lastly, (vi) melt-quenched amorphous.

enable rapid model convergence in the Σ_{AI} loop while remaining flexible with respect to the gradient map in synthesis phase space.

In contrast to the X_{AI} cycle, there is more opportunity to incorporate structure based on prior knowledge into the acquisition function, rather than the kernel, of Σ_{AI}. As shown in Fig. 3C, random sampling performs only slightly worse than more sophisticated acquisition methods like uncertainty sampling or upper confidence bound (UCB) sampling in terms of R_s^2, a generalization of R^2 that takes into account the heteroscedasticity of the data due to the propagation of uncertainty (see the "Error metrics" section). This behavior can be understood by considering the following: Every experiment at $\{T_p, \tau\}$ produces a range of temperatures $T < T_p$ at which new information is obtained, thereby reducing the uncertainty not only at T_p but also in a wide range of temperatures below it. Hence, uncertainty sampling at T_p and τ alone is a poor strategy. To address this issue, we introduce stripe uncertainty (SU) sampling, which takes into account the uncertainty in the whole temperature range between T_{min} and T_p. This strategy greatly improves performance, reaching $R_s^2 > 0.7$ within 15 iterations.

The plot in Fig. 3D shows the gradient heatmap of Bi_2O_3 from an exhaustive sampling of all 617 lg-LSA stripes, with gradient peaks at every value of τ highlighted in white. Note that these peaks are connected and form ridges that can be well interpreted as phase field boundaries. To label these phase fields, we selectively collect and analyze XRD data of lg-LSA stripes annealed at conditions close to the field centers (see the Supplementary Materials for details). Notably, only few XRD measurements suffice to label the phase map, once the phase boundaries have been determined via the reflectance data. The phase field (i) below approximately 350°C corresponds to the as-deposited amorphous film, while a slight gradient ridge separates it from (ii), a densified, amorphous

regime. At approximately 500°C, there is a ridge that extends across the complete dwell range, corresponding to the crystallization onset of the δ phase in domain (iii). The boundary separating (iii) from (iv) at approximately 550°C corresponds to the onset of a two-phase region, where both the δ and β phases of Bi_2O_3 coexist, and above approximately 650°C, we observe phase-pure β-Bi_2O_3 in (v). The gradient ridge between (iv) and (v) is particularly weak and wide in T, and we consider the bump arising at around $10^{3.5}$ μs to be an artifact within the measurement uncertainty. The phase field above approximately 810°C corresponds to amorphous Bi_2O_3 that reforms after quenching from melt [the bulk melting temperature of Bi_2O_3 is 817°C (50)] and stretches out across all values of τ.

A representative evolution of the actively learned gradient phase map is shown in Fig. 4, with six snapshots from (A) to (F). Figure 4A shows the preliminary gradient map at iteration $n = 3$: The two gradient ridges spanning all dwells qualitatively correspond to the crystallization boundaries from either melt (v and vi in Fig. 3D) or the deposited, densified thin film (ii and iii). At $n = 8$ in Fig. 4B, we detect the onset of the two-phase region (iii and iv), and at $n = 15$ (Fig. 4C), we detect the phase-pure β-Bi_2O_3 boundary (iv and v). With only $n = 25$ iterations, we identify the last boundary, namely, the amorphous densification onset (i and ii; see Fig. 4D). At this point, the overall features of the gradient phase map are already qualitatively captured completely, and subsequent iterations merely refine the boundary locations (Fig. 4E), getting closer to the exhaustive phase map in Fig. 4F.

Two factors are crucial for Σ_{AI} to achieve a factor of approximately 14 acceleration to reach $R_s^2 = 0.7$ compared to random sampling without propagation of input uncertainties. First, incorporating materials synthesis into our SARA discovery framework allows us to check for convergence of the phase diagram on the fly. Even with random sampling, the possibility of quantifying the progress and
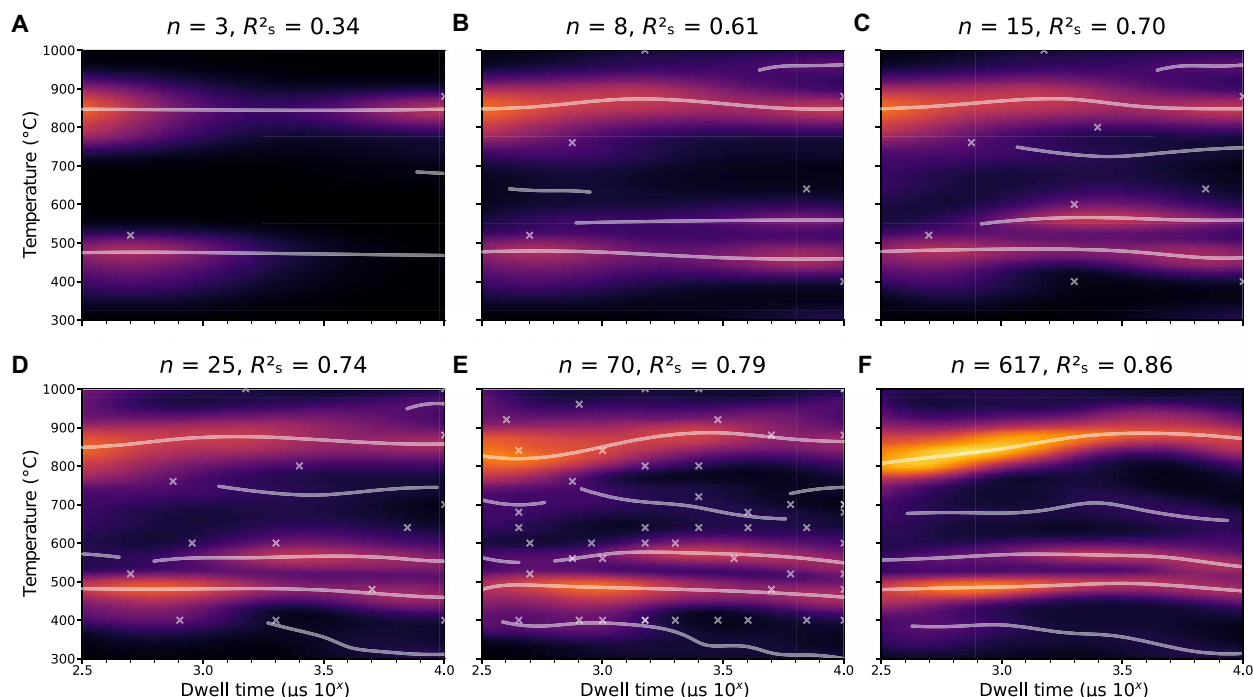
Fig. 4. The evolution of the actively learned gradient phase map of the Bi$_2$O$_3$ system at selected number of iterations n. (**A** to **E**) We use the stripe uncertainty (SU) acquisition strategy, starting from a randomly selected condition $(T_p, \tau)_1$. The gradient ridges are shown as white lines, and the conditions $(T_p, \tau)_i$ at which the experiments have been performed are shown as white crosses (note that not all crosses are shown, because the plots have been cropped to a smaller range than the range of experimental conditions). For each panel, the number of sampled conditions n is indicated at the top, together with the corresponding R_s^2 score. (**F**) Final exhaustively sampled phase map.

monitoring convergence in the gradient mapping informs us how well the phase space has been sampled, thereby substantially reducing the resource cost. Second, the comprehensive uncertainty propagation in conjunction with the stripe uncertainty acquisition function realizes the full potential of AI and AL and decreases the required samples to a fraction of the exhaustive measurements. SARA's overall AL acceleration is the product of the acceleration factors of X_{AI} and Σ_{AI} due to the cycles' nested design.

DISCUSSION
In conclusion, we have developed SARA, an AI-driven autonomous closed-loop materials discovery framework that integrates robotic materials synthesis with automated microscopy imaging and reflectance spectroscopy characterization. SARA incorporates a set of nested AL loops based on specialized physics-inspired GP regression models to synthesize, characterize, and iteratively explore nonequilibrium synthesis phase maps using HT lg-LSA thin-film processing. In particular, SARA tightly integrates the physics of the experiments and quantifies experimental uncertainties in both the inputs and the outputs of the model. We highlight SARA's capabilities on the Bi$_2$O$_3$ system by showing that SARA reduces the time to map the system's phase boundaries by more than two orders of magnitude, in contrast to random or exhaustive searches. In particular, SARA identifies the synthesis conditions that trap metastable δ-Bi$_2$O$_3$ at room temperature, a promising solid oxide electrolyte. While scaling-up synthesis is a challenge for future work, a flat-top laser profile could be applied to anneal Bi$_2$O$_3$ at these processing conditions and

integrated into fabrication of thin-film solid oxide fuel cells in portable power applications, or other integrated solid-state micro-electromechanical system (MEMS) devices (*51*, *52*).

The speedup in synthesis and data acquisition achieved by SARA is a fundamental prerequisite for paving the path toward exploratory HT efforts with additional chemical degrees of freedom and extended processing parameters and when targeting property optimization. The gradient phase map construction can be extended to additional degrees of freedom, e.g., on composition spreads over a continuous range of chemistries. While our current lg-LSA synthesis is limited to inorganic thin films that are transparent to infrared (IR) laser radiation (*33*), e.g., complex oxides, future efforts are aimed at incorporating IR-transparent substrates that will allow the autonomous synthesis and processing of broader materials classes, e.g., metals and alloys. Furthermore, techniques that enable the characterization of *in situ* lg-LSA processing would provide us the means to better understand the transformation kinetics of metastable phases and improve our design of physics-inspired AI models. SARA's nested AI architecture also allows the incorporation of additional agents for multi-objective optimization efforts by including robotic measurements of target properties. In addition to phase boundary mapping, research objectives for which SARA would enable new modalities of materials design include the following: discovery of a synthesis condition for a not-yet-synthesized phase, extension of the optical spectroscopy to characterize visible absorption to identify syntheses of materials for solar energy applications, and incorporation of new performance characterization such as electrical conductivity measurements. These latter examples involve

mapping of synthesis phase diagrams in the context of performance metrics for a target application, the central goal of studying PCSP relationships. SARA's autonomous execution of these studies constitutes a grand vision of AI-assisted materials science.

MATERIALS AND METHODS
Experiments and measurements
Thin-film deposition
We used thermally oxidized (200 nm oxide), highly doped (p-type, 0.01 to 0.02 Ohm·cm) Si wafers with lithographically patterned gold alignment marks as substrates for our thin-film deposition. Radio frequency (RF) reactive sputtering from a Bi target in an atmosphere of 8-mTorr Ar and 2-mTorr O_2 was used to deposit the Bi_2O_3 sample in a custom-built sputter system. The substrate was rotated while operating the target at an RF power of 20 W to create a 170-nm-thick film with <10% thickness variation.

Lateral gradient laser spike annealing
The lg-LSA anneals were conducted using a continuous-wave CO_2 laser operating at $\lambda = 10.6$ μm and maximum power of 125 W, which was configured to produce a power profile with a bi-Gaussian shape (320-μm-wide lateral FWHM and 80-μm-long longitudinal FWHM). To reach steady state, each anneal was conducted on a 5-mm-long stripe, with peak temperatures ranging from 400° to 1300°C and processing dwell times between 250 μs and 10 ms. The stripes were located 2 mm apart to avoid thermal overlap between anneals. With this configuration, a 100-mm-diameter wafer offers space for a total of up to 625 stripes with distinct anneal conditions. Note that the dwell τ is related to the scan velocity v via the FWHM of the laser in the scan direction (longitudinal) through $\tau = \frac{\text{FWHM}}{v}$. τ is approximately the time scale during which the temperature is within 5% of the peak temperature (33). To avoid potential location bias on the wafer arising from variations in film thickness, the anneal locations were randomized across the thin film with respect to T_p and τ. In total, we annealed 617 lg-LSA stripes on our Bi_2O_3 sample with $400 \leq T_p \leq 1300°$C and $250 \leq T_p \leq 10{,}000$ μs.

Microscopy imaging
We used a Thorlabs complementary metal-oxide semiconductor camera (RGB channels with 1024×1280 pixels), which was aligned normal to the sample, together with a coaxial illumination using white light over a spot size of approximately 1 mm in diameter. The camera magnification was set to produce a field of view of approximately 1 mm horizontally, resulting in a spacing of 0.92 μm between pixels. The raw microscope images and the software to process them are available online (53, 54).

Reflectance spectroscopy
A white light source ($400 < \lambda < 900$ nm) was focused down to a single 10-μm-diameter spot using optical fibers to locally illuminate the sample, allowing spatially resolved reflectance measurements. We used a flame spectrometer from Ocean Optics to collect the reflectance spectroscopy with an optimized integration time of ≈4500 ms. The reflected light was measured from $\lambda = 340$ to $\lambda < 1026$ nm at 2046 discrete values. The reflectance data were calibrated and normalized with respect to a dark reference spectrum, and a spectrum from an Ag-coated mirror. For the exhaustive reflectance measurements, the optical fiber was scanned across an lg-LSA stripe over a range of 1.5 mm in 10-μm increments, leading to 151 samples per stripe. The raw reflectance data and the software to process are available online (53, 54).

X-ray diffraction
The XRD data were collected using the ID3B beamline at the Cornell High Energy Synchrotron Source (CHESS) with a 9.7-keV beam, which was focused to a spot size on the sample of 20 μm by 40 μm at a 2° angle of incidence. A Pilatus 300K detector was used to capture the diffracted signal. The XRD data were collected every 10 μm across a stripe with a 50-ms integration time for each frame. The 2D detector data were integrated along the χ direction using pyFAI (55).

Computational methods
In the following, bold lowercase letters refer to vectors and bold uppercase letters refer to matrices. Given a collection of inputs $\mathbf{X} = [\mathbf{x}_1, …, \mathbf{x}_n]$ of a function f, we let $f(\mathbf{X})$ be the result of the application of f to each column of $\mathbf{X}$, $f(\mathbf{X}) \coloneqq [f(\mathbf{x}_1), …, f(\mathbf{x}_N)]$.

Gaussian processes
A GP is a distribution over functions whose finite-dimensional marginal distributions are multivariate normal. That is, for any sample f of a GP and any finite selection of inputs $\mathbf{X}$, we have $f(\mathbf{X}) \sim \mathcal{N}(\boldsymbol{\mu_X}, \Sigma_{\mathbf{X}})$, for some mean vector $\boldsymbol{\mu_X}$ and covariance matrix $\Sigma_{\mathbf{X}}$. Analogous to the multivariate case, a GP is completely defined by its first and second moments: a mean function $\mu(\cdot)$ and a covariance function $\kappa(\cdot,\cdot)$, also known as a kernel. In particular, if $f \sim \mathcal{GP}(\mu, \kappa)$, then for any finite collection of inputs $\mathbf{X}$

$$f(\mathbf{X}) \sim \mathcal{N}(\mu(\mathbf{X}), \kappa(\mathbf{X}, \mathbf{X})) \tag{1}$$

where $\kappa(\mathbf{X}, \mathbf{X})$ is the matrix whose $(i, j)^{\text{nth}}$ entry is $\kappa(\mathbf{x}_i, \mathbf{x}_j)$. Fortunately, the posterior mean μ_p and posterior covariance κ_p of a GP conditioned on observations with normally distributed noise have closed forms and only require linear algebraic operations

$$
\begin{aligned}
\mu_p(\mathbf{x}_*) &= \mu(\mathbf{x}_*) + \kappa(\mathbf{x}_*, \mathbf{X})\, \Sigma_{\mathbf{X}}^{-1}(\mathbf{y} - \mu(\mathbf{X})) \\
\kappa_p(\mathbf{x}_*, \mathbf{x}'_*) &= \kappa(\mathbf{x}_*, \mathbf{x}'_*) - \kappa(\mathbf{x}_*, \mathbf{X})\, \Sigma_{\mathbf{X}}^{-1} \kappa(\mathbf{X}, \mathbf{x}'_*)
\end{aligned} \tag{2}
$$

where for homoscedastic regression $\Sigma_{\mathbf{X}} = \kappa(\mathbf{X}, \mathbf{X}) + \sigma_y^2 \mathbf{I}$ and σ_y is the SE of the target $\mathbf{y}$. Because a GP's behavior is chiefly determined by the kernel, its performance can be improved markedly by incorporating important problem structure into the kernel. For more background on GPs, see (56). For the present work, we developed a GP framework in Julia (57) with which we implemented SARA's AL technology (54).

Active learning
The field of AL considers the problem of selecting data in an optimal way to reduce the total amount of data that is required to effectively train a model (5, 58). To this end, the notion of an acquisition function is important. An acquisition function $a(\mathbf{X}, \mathbf{y})$ depends on currently available data and outputs a suggested observation $\mathbf{x}^*$. For example, if $f\,|\,\mathbf{X}, \mathbf{y}$ denotes the posterior of f after having seen the data, and $\text{var}(f\,|\,\mathbf{X}, \mathbf{y})$ is the posterior variance (itself a function), then

$$\underset{\mathbf{x}^*}{\arg\max}\, \text{var}(f\,|\,\mathbf{X}, \mathbf{y})\,(\mathbf{z}) \tag{3}$$

defines an acquisition strategy known as uncertainty sampling. Other acquisition functions are based on upper-confidence bounds, expected improvement, and probability of improvement. Overall, an important ingredient for AL is the quantification of uncertainty,

which is a strength of Bayesian models. In the realm of Bayesian models, GPs are of particular importance because of their unique combination of flexibility, closed-form inference formulas, and uncertainty quantification. For these reasons, we chose to build SARA's computational backbone on GPs.

Input noise

Because of the importance of uncertainty quantification for AL, it is critical to take all sources of uncertainty into account. In the case of SARA, it is crucial to account for errors in not only the measurements (i.e., model outputs) but also the experimental conditions (i.e., model inputs) due to intrinsic experimental uncertainties in the temperature profile. However, the general problem of posterior inference with input noise is intractable. For this reason, one needs to use approximate methods like variational approximations (59, 60) and Markov chain Monte Carlo (61) or methods that transform the problem of homoscedastic regression with input noise to one of heteroscedastic regression (62) without input noise (63–65). A particularly efficient technique is that of McHutchon and Rasmussen (66), which is based on propagating the input uncertainty using a linear approximation of the standard posterior mean. According to this model, given the regular posterior mean $\mu_p(x)$, the input noise–corrected version can be computed by updating

$$\Sigma_{\mathbf{X}} \leftarrow \Sigma_{\mathbf{X}} + \mathrm{diag}\left(\sigma_x(\mathbf{X}) \odot \partial_x \mu_p(\mathbf{X})\right)^2 \tag{4}$$

in Eq. 2 for the GP posterior. Notably, we generalize the original work in making the input uncertainty $\sigma_x(\mathbf{X})$ dependent on the input. This is possible because the non-constant uncertainties in SARA's experimental process can be estimated well by physical considerations (see the "Σ_{AI}" section for details). Lastly, note that Eq. 4 makes the approximate posterior uncertainty dependent on the values of the data via the posterior mean, not just the locations of the measurements.

X_{AI}

The goal of X_{AI} is to infer the reflectance $r(x, \lambda)$ using the least number of measurement locations x_i as possible. Each measurement of the inner loop acquires the wavelength-dependent spectroscopic reflectance of the underlying thin film, that is, a vector whose entries correspond to reflectance intensities at a given wavelength.

To aid the efficiency of our model, we first reduce the dimension of the output by projecting it onto the basis of a small number (10 to 20) of Legendre polynomials. Because the signal is smooth as a function of wavelength, it admits a sparse approximation in this basis, allowing the compression of the signal with virtually no loss of information (49) [see also the Supplementary Materials]. The AL cycle then works on the dimensionality-reduced form of the reflectance data.

In the following, we describe the construction of the X_{AI} kernel function, which integrates special structure of the data and is a critical part of X_{AI}. In particular, the kernel incorporates (i) lateral symmetry, (ii) variance scaling based on RGB data, and (iii) asymptotically linear behavior. Starting with a Matérn 5/2 kernel k with a length scale l, we symmetrize it via $k_{sym}(x, y) = k(x, y) + k(x - c, y - c)$ around the stripe center c, which we estimate from the RGB images. We incorporate further information from the RGB images by scaling the kernel with the LSA or RGB prior function f_{rgb} shown in Fig. 2C. In particular, we use the peaks in the RGB gradient signal, slightly broaden them by a Gaussian with $\sigma = 20\ \mu m$, and sum them to our RGB prior function (purple line in Fig. 2C). In addition, the

overall width of the lg-LSA stripe gives rise to the LSA prior, which is a generalized Gaussian with a wide shape parameter of $\beta = 4$ and a scale parameter σ defined by the stripe width (red line in Fig. 2C). f_{rgb} is then given by a weighted sum of these two prior functions. This scaling constrains the search space, because we do not expect a lot of change in the underlying material if the experimental conditions (e.g., temperature) stay similar, and gives rise to the kernel $f_{rgb}(x)k_{sym}(x, y)f_{rgb}(y)$. Lastly, we incorporate an asymptotically linear behavior, due to thickness variations in the wafer, with the linear kernel $k_{line}(x, y) = x \cdot y + b$, where b is a constant that controls the variance of the bias term of the line. As a result, the X_{AI} kernel for one Legendre coefficient is proportional to

$$k_{X_{AI}}(x, y) = f_{rgb}(x)\,k_{sym}(x, y)\,f_{rgb}(y) + k_{line}(x, y) \tag{5}$$

For all the Legendre coefficients, we then use a GP with the kernel $a_i k_{X_{AI}}(x, y)$, where $\{a_i\}$ are scaling coefficients that incorporate the different variances of the Legendre coefficients, to learn the reflectance map. This can also be interpreted as computing a vector-valued GP with the matrix-valued kernel $\mathbf{K}_{X_{AI}}(x, y) = \mathrm{diag}\,(\mathbf{a})\,k_{X_{AI}}(x, y)$, where $\mathbf{a}$ is the vector of scaling coefficients. For a comprehensive review on matrix-valued kernels, see (67). The length scale l of the Matérn kernel can be optimized via maximization of the marginal likelihood (56). However, to make the reported results in Fig. 2D independent of this nonconvex optimization procedure, we ran the benchmarks using a range of fixed length scales and reported the best performing combination for each kernel.

Regarding the acquisition function, in addition to uncertainty sampling, we benchmark X_{AI} using IU sampling, a policy that reduces the total variance over a set of potential measurement locations $\mathbf{Z}$. In particular, IU is defined by

$$\arg\min_{\mathbf{x}^*} \sum_{\mathbf{z} \in \mathbf{Z}} \mathrm{var}(f \mid \mathbf{X}, \mathbf{y}, \mathbf{x}^*)(\mathbf{z}) \tag{6}$$

where $\mathbf{X}$ is the set of inputs and $\mathbf{y}$ is the set of outputs of the model. Note that we can calculate the quantity $\mathrm{var}(f \mid \mathbf{X}, \mathbf{y}, \mathbf{x}^*)$ because the standard posterior GP variance only depends on the measurement location, not the value y^*. Lastly, we note that the derivative of a GP is also a GP (68). Plugging the derivative GP into Eq. 6 yields IGU sampling, which achieves the best performance in the X_{AI} acquisition benchmark (see Fig. 2).

Σ_{AI}

Σ_{AI} works to identify phase regions and their boundaries in the temperature–dwell time space and, more generally, the processing-composition space. The raw reflectance data cannot be used directly for this task because of two reasons. First, the data are measured as a function of position, not temperature. Therefore, we convert the stripe-specific reflectance function $r_{T_p,\tau}(x, \lambda)$ to the temperature domain using the temperature profile $T_{T_p,\tau}$, yielding $r_{T_p,\tau}(T, \lambda)$. Second, the reflectance varies not only with the phase behavior but also with the film thickness across the wafer. For this reason, we calculate the L_2-norm of the rate of change of the spectroscopic reflectance, which is invariant to linear thickness variations of the film. In particular, for all $T < T_p$, we want to infer

$$d(T, \tau) := \sqrt{\int \left(\frac{\partial\, r_{T_p,\tau}(T, \lambda)}{\partial T}\right)^2 \mathrm{d}\lambda} \tag{7}$$

d quantifies how much the spectroscopic reflectance changes as a function of temperature and dwell time and is a strong indicator of phase changes (*47*). Estimating the phase boundaries then reduces to getting an accurate estimate of d over all (T, τ) (and potentially composition c). This is the goal of the Σ_{AI} loop.

Crucially, experimental errors can occur in x and, therefore, in T, making it imperative to quantify the uncertainty due to these input errors and propagate them to Σ_{AI}. Our benchmarks show that ignoring these uncertainties leads to a substantial deterioration in AL performance (see Fig. 3C).

To this end, we now discuss the intrinsic experimental uncertainties due to the temperature profile $T_{T_p,\tau}(x)$ of the laser. In particular, we compute the variance of the true temperature around the value predicted by the temperature profile as a function of position by

$$\sigma_T^2(x) = \sigma_{T_p}^2 \left(\frac{T_p}{1400} \right)^2 + \sigma_x^2 \left(\frac{\partial T_{T_p,\tau}(x)}{\partial x} \right)^2 \qquad (8)$$

where σ_{T_p} is the SE in the peak temperature and σ_x is the SE in the position. The first term quantifies the error at the peak temperature, which is largest at high temperatures (1400°C) and falls off linearly with T. The second term quantifies uncertainties of the temperature profile, which not only are due to limited spatial resolution but also encompass random asymmetries in the profile of the laser. The form of term is derived using SE propagation techniques (*69*). For the results reported here, $\sigma_{T_p} = 25$°C and $\sigma_x = 50$ μm.

The expression for the temperature uncertainty in Eq. 8 is then used in conjunction with Eq. 4 to compute a GP that comprehensively quantifies the uncertainties in the Legendre coefficients of the optical reflectance as a function of temperature. To compute d in Eq. 7, one simply sums the squared derivatives of the GPs of the Legendre coefficients of the reflectance

$$d(T, \tau) = \sqrt{\sum_i \left(\partial_T \mu_p^{(i)}(T) \right)^2} \qquad (9)$$

Because we have access to the uncertainties in $\mu_p^{(i)}$ from the GP, we can use uncertainty propagation techniques on Eq. 9 to calculate a first-order uncertainty estimate of $d(T, \tau)$. For the outer loop, we used a two-dimensional Matérn 5/2 kernel with different length scales across each dimension. This allows the GP to learn independent sensitivity parameters of the experiment for the input dimensions. Note that the Σ_{AI} benchmarks in Fig. 3C were carried out with fixed length scales to disentangle the effects of different acquisition functions and hyperparameter learning.

For Σ_{AI}, we designed an acquisition strategy that incorporates the property that a single stripe generates data throughout a range of temperatures. In particular, given experimental conditions $\mathbf{x}_s$ that give rise to a stripe (T_p, τ, etc.), we sum the uncertainties of all relevant observations $\mathbf{x}_i$ that are in the set Stripe($\mathbf{x}_s$) of conditions on the stripe $\mathbf{x}_s$. In particular, we propose stripe uncertainty sampling:

$$\arg\max_{\mathbf{x}_s} \sum_{\mathbf{x}_i \in \text{Stripe}(\mathbf{x}_s)} \text{var}(f|X, y)(x_i) \qquad (10)$$

Notably, one can use the same principle to generalize other acquisition functions. We investigated a stripe upper-confidence bound sampling policy. However, it performed worse or equal to the simpler stripe uncertainty sampling policy above. The synergy of the comprehensive uncertainty quantification and the stripe sampling function yields considerable benefits, as displayed in Fig. 3C.

Error metrics

In our benchmarks of the kernels and acquisition functions for the inner loop, we used the coefficient of determination R^2 to measure performance, defined by

$$R^2 = 1 - \frac{\sum_i (f(x_i) - y_i)^2}{\sum_i (\mu(\mathbf{y}) - y_i)^2} \qquad (11)$$

where $\mu(\mathbf{y})$ is the mean of the data $\mathbf{y}$. The advantage of using R^2 over other canonical measures like the mean-squared error is that it is dimensionless and easily interpretable as the proportion of the variance of the data that is explained by the model f.

As R^2 weighs the deviation at every data point equally, it is not an ideal measure for heteroscedastic data, like the optical gradient data of Σ_{AI}. For this reason, we use a generalization of R^2, based on the log-likelihood of the heteroscedastic normal errors, to measure performance in the Σ_{AI} benchmarks. In particular, the measure is given by

$$R_s^2 = 1 - \frac{\sum_i (f(x_i) - y_i)^2 / \sigma_i^2}{\sum_i (\mu(\mathbf{y}) - y_i)^2 / \sigma_i^2} \qquad (12)$$

where σ_i is the SD of the i^{th} error. For SARA, the σ_i are the product of the comprehensive uncertainty quantification of the experimental process. Clearly, R_s^2 reduces to R^2 if the noise variances are all equal. If they are not, R_s^2 is a better measure of misfit, as it weighs the residuals of more certain data points stronger than those with greater uncertainty. Notably, similar pseudo-R^2 scores based on log-likelihoods are used throughout statistics and applied fields (*70–72*).

SUPPLEMENTARY MATERIALS

REFERENCES AND NOTES

1. A. W. Senior, R. Evans, J. Jumper, J. Kirkpatrick, L. Sifre, T. Green, C. Qin, A. Žídek, A. W. R. Nelson, A. Bridgland, H. Penedones, S. Petersen, K. Simonyan, S. Crossan, P. Kohli, D. T. Jones, D. Silver, K. Kavukcuoglu, D. Hassabis, Improved protein structure prediction using potentials from deep learning. *Nature* **577**, 706–710 (2020).
2. B. Huang, O. A. von Lilienfeld, Quantum machine learning using atom-in-molecule-based fragments selected on the fly. *Nat. Chem.* **12**, 945–951 (2020).
3. G. Carleo, M. Troyer, Solving the quantum many-body problem with artificial neural networks. *Science* **355**, 602–606 (2017).
4. N. Jean, M. Burke, M. Xie, W. M. Davis, D. B. Lobell, S. Ermon, Combining satellite imagery and machine learning to predict poverty. *Science* **353**, 790–794 (2016).
5. B. Settles, "Active learning literature survey" (Technical reports, University of Wisconsin-Madison Department of Computer Sciences, 2009).
6. Y. Gil, M. Greaves, J. Hendler, H. Hirsh, Amplify scientific discovery with artificial intelligence. *Science* **346**, 171–172 (2014).
7. A. A. Melnikov, H. Poulsen Nautrup, M. Krenn, V. Dunjko, M. Tiersch, A. Zeilinger, H. J. Briegel, Active learning machine learns to create new quantum experiments. *Proc. Natl. Acad. Sci. U.S.A.* **115**, 1221–1226 (2018).
8. R. F. Murphy, An active role for machine learning in drug development. *Nat. Chem. Biol.* **7**, 327–330 (2011).
9. J. Z. Kolter, Z. Jackowski, R. Tedrake, Design, analysis, and learning control of a fully actuated micro wind turbine, in *Proceedings of the 2012 American Control Conference (ACC)* (IEEE, 2012), pp. 2256–2263.
10. R. B. van Dover, L. F. Schneemeyer, The codeposited composition spread approach to high-throughput discovery/exploration of inorganic materials. *Macromol. Rapid Commun.* **25**, 150–157 (2004).

11. J. R. Hattrick-Simpers, J. M. Gregoire, A. G. Kusne, Perspective: Composition-structure-property mapping in high-throughput experiments: Turning data into knowledge. *APL Materials* **4**, 053211 (2016).

12. K. G. Reyes, B. Maruyama, The machine learning revolution in materials? *MRS Bull.* **44**, 530–537 (2019).

13. H. S. Stein, J. M. Gregoire, Progress and prospects for accelerating materials science with automated and autonomous workflows. *Chem. Sci.* **10**, 9640–9649 (2019).

14. 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, A. Aspuru-Guzik, Accelerating the discovery of materials for clean energy in the era of smart automation. *Nat. Rev. Mater.* **3**, 5–20 (2018).

15. 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, I. Takeuchi, On-the-fly closed-loop materials discovery via Bayesian active learning. *Nat. Commun.* **11**, 5966 (2020).

16. D. Xue, P. V. Balachandran, J. Hogden, J. Theiler, D. Xue, T. Lookman, Accelerated search for materials with targeted properties by adaptive design. *Nat. Commun.* **7**, 11241 (2016).

17. R. Yuan, Z. Liu, P. V. Balachandran, D. Xue, Y. Zhou, X. Ding, J. Sun, D. Xue, T. Lookman, Accelerated discovery of large electrostrains in BaTiO3-based piezoelectrics using active learning. *Adv. Mater.* **30**, 1702884 (2018).

18. P. Nikolaev, D. Hooper, N. Perea-López, M. Terrones, B. Maruyama, Discovery of wall-selective carbon nanotube growth conditions via automated experimentation. *ACS Nano* **8**, 10214–10222 (2014).

19. P. Nikolaev, D. Hooper, F. Webber, R. Rao, K. Decker, M. Krein, J. Poleski, R. Barto, B. Maruyama, Autonomy in materials research: A case study in carbon nanotube growth. *npj Comput. Mater.* **2**, 16031 (2016).

20. S. Sun, N. T. P. Hartono, Z. D. Ren, F. Oviedo, A. M. Buscemi, M. Layurova, D. X. Chen, T. Ogunfunmi, J. Thapa, S. Ramasamy, C. Settens, B. L. DeCost, A. G. Kusne, Z. Liu, S. I. P. Tian, I. M. Peters, J.-P. Correa-Baena, T. Buonassisi, Accelerated development of perovskite-inspired materials via high-throughput synthesis and machine-learning diagnosis. *Joule* **3**, 1437–1451 (2019).

21. R. W. Epps, M. S. Bowen, A. A. Volk, K. Abdel-Latif, S. Han, K. G. Reyes, A. Amassian, M. Abolhasani, Artificial chemist: An autonomous quantum dot synthesis bot. *Adv. Mater.* **32**, 2001626 (2020).

22. B. P. MacLeod, F. G. L. Parlane, T. D. Morrissey, F. Häse, L. M. Roch, K. E. Dettelbach, R. Moreira, L. P. E. Yunker, M. B. Rooney, J. R. Deeth, V. Lai, G. J. Ng, H. Situ, R. H. Zhang, M. S. Elliott, T. H. Haley, D. J. Dvorak, A. Aspuru-Guzik, J. E. Hein, C. P. Berlinguette, Self-driving laboratory for accelerated discovery of thin-film materials. *Sci. Adv.* **6**, eaaz8867 (2020).

23. A. E. Gongora, B. Xu, W. Perry, C. Okoye, P. Riley, K. G. Reyes, E. F. Morgan, K. A. Brown, A Bayesian experimental autonomous researcher for mechanical design. *Sci. Adv.* **6**, eaaz1708 (2020).

24. L. Porwol, D. J. Kowalski, A. Henson, D.-L. Long, N. L. Bell, L. Cronin, An autonomous chemical robot discovers the rules of inorganic coordination chemistry without prior knowledge. *Angew. Chem. Int. Ed.* **59**, 11256–11261 (2020).

25. J. Li, J. Li, R. Liu, Y. Tu, Y. Li, J. Cheng, T. He, X. Zhu, Autonomous discovery of optically active chiral inorganic perovskite nanocrystals through an intelligent cloud lab. *Nat. Commun.* **11**, 2046 (2020).

26. Z. Li, M. A. Najeeb, L. Alves, A. Z. Sherman, V. Shekar, P. Cruz Parrilla, I. M. Pendleton, W. Wang, P. W. Nega, M. Zeller, J. Schrier, A. J. Norquist, E. M. Chan, Robot-accelerated perovskite investigation and discovery. *Chem. Mater.* **32**, 5650–5663 (2020).

27. K. Alberi, M. B. Nardelli, A. Zakutayev, L. Mitas, S. Curtarolo, A. Jain, M. Fornari, N. Marzari, I. Takeuchi, M. L. Green, M. Kanatzidis, M. F. Toney, S. Butenko, B. Meredig, S. Lany, U. Kattner, A. Davydov, E. S. Toberer, V. Stevanovic, A. Walsh, N.-G. Park, A. Aspuru-Guzik, D. P. Tabor, J. Nelson, J. Murphy, A. Setlur, J. Gregoire, H. Li, R. Xiao, A. Ludwig, L. W. Martin, A. M. Rappe, S.-H. Wei, J. Perkins, The 2019 materials by design roadmap. *J. Phys. D. Appl. Phys.* **52**, 013001 (2018).

28. A. Saksena, Y.-C. Chien, K. Chang, P. Kümmerl, M. Hans, B. Völker, J. M. Schneider, Metastable phase formation of Pt-X (X = Ir, Au) thin films. *Sci. Rep.* **8**, 10198 (2018).

29. L. Li, K. Jamieson, G. DeSalvo, A. Rostamizadeh, A. Talwalkar, Hyperband: A novel bandit-based approach to hyperparameter optimization. *J. Mach. Learn. Res.* **18**, 6765–6816 (2017).

30. M. O. Ahmed, B. Shahriari, M. Schmidt, Do we need "harmless" Bayesian optimization and "first-order" Bayesian optimization. *NIPS BayesOpt* (2016).

31. Scientific Autonomous Reasoning Agent (SARA): Integrating Materials Theory, Experiment and Computation—Research Areas—AFOSR.

32. P. V. Balachandran, B. Kowalski, A. Sehirlioglu, T. Lookman, Experimental search for high-temperature ferroelectric perovskites guided by two-step machine learning. *Nat. Commun.* **9**, 1668 (2018).

33. R. T. Bell, A. G. Jacobs, V. C. Sorg, B. Jung, M. O. Hill, B. E. Treml, M. O. Thompson, Lateral temperature-gradient method for high-throughput characterization of material processing by millisecond laser annealing. *ACS Comb. Sci.* **18**, 548–558 (2016).

34. T. F. T. Cerqueira, R. Sarmiento-Pérez, M. Amsler, F. Nogueira, S. Botti, M. A. L. Marques, Materials design on-the-fly. *J. Chem. Theory Comput.* **11**, 3955–3960 (2015).

35. S. Kirklin, J. E. Saal, B. Meredig, A. Thompson, J. W. Doak, M. Aykol, S. Rühl, C. Wolverton, The Open Quantum Materials Database (OQMD): Assessing the accuracy of DFT formation energies. *npj Comput. Mater.* **1**, 15010 (2015).

36. D. R. Clarke, S. R. Phillpot, Thermal barrier coating materials. *Mater. Today* **8**, 22–29 (2005).

37. R. H. J. Hannink, P. M. Kelly, B. C. Muddle, Transformation toughening in zirconia-containing ceramics. *J. Am. Ceram. Soc.* **83**, 461–487 (2000).

38. D. R. Clarke, C. G. Levi, Materials design for the next generation thermal barrier coatings. *Annu. Rev. Mater. Res.* **33**, 383–417 (2003).

39. N. Satoh, T. Nakashima, K. Yamamoto, Metastability of anatase: Size dependent and irreversible anatase-rutile phase transition in atomic-level precise titania. *Sci. Rep.* **3**, 1959 (2013).

40. Z.-H. Cui, F. Wu, H. Jiang, First-principles study of relative stability of rutile and anatase TiO2 using the random phase approximation. *Phys. Chem. Chem. Phys.* **18**, 29914–29922 (2016).

41. N. H. Vu, H. V. Le, T. M. Cao, V. V. Pham, H. M. Le, D. Nguyen-Manh, Anatase-rutile phase transformation of titanium dioxide bulk material: A DFT + U approach. *J. Phys. Condens. Matter* **24**, 405501 (2012).

42. H. A. Harwig, On the structure of bismuthsesquioxide: The α, β, γ, and δ-phase. *Z. Anorg. Allg. Chem.* **444**, 151–166 (1978).

43. N. Cornei, N. Tancret, F. Abraham, O. Mentré, New ε-Bi2O3 metastable polymorph. *Inorg. Chem.* **45**, 4886–4888 (2006).

44. P. Shuk, H. D. Wiemhöfer, U. Guth, W. Göpel, M. Greenblatt, Oxide ion conducting solid electrolytes based on Bi2O3. *Solid State Ionics* **89**, 179–196 (1996).

45. R. T. Bell, P. A. Beaucage, M. J. Murphy, A. B. Connolly, U. Wiesner, R. B. Van Dover, 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**, 4328–4336 (2021).

46. V. Borisenko, P. Hesketh, *Rapid Thermal Processing of Semiconductors* (Springer, 1997).

47. 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, M. O. Thompson, Optical identification of materials transformations in oxide thin films. *ACS Comb. Sci.* **22**, 887–894 (2020).

48. S. Thienhaus, D. Naujoks, J. Pfetzing-Micklich, D. König, A. Ludwig, Rapid identification of areas of interest in thin film materials libraries by combining electrical, optical, x-ray diffraction, and mechanical high-throughput measurements: A case study for the system Ni–Al. *ACS Comb. Sci.* **16**, 686–694 (2014).

49. H. Wang, S. Xiang, On the convergence rates of Legendre approximation. *Math. Comput.* **81**, 861 (2012).

50. O. Madelung, U. Rössler, M. Schulz, Eds., SpringerMaterials. Bismuth oxide (Bi2O3) crystal structure, chemical bond, lattice parameters. Landolt-Börnstein - Group III Condensed Matter 41C (*Non-Tetrahedrally Bonded Elements and Binary Compounds I*).

51. A. Evans, A. Bieberle-Hütter, J. L. Rupp, L. J. Gauckler, Review on microfabricated micro-solid oxide fuel cell membranes. *J. Power Sources* **194**, 119–129 (2009).

52. M. Tsuchiya, B.-K. Lai, S. Ramanathan, Scalable nanostructured membranes for solid-oxide fuel cells. *Nat. Nanotech.* **6**, 282–286 (2011).

53. Raw microscope images and reflectance data for the lg-LSA annealed Bi2O3 thin film; https://doi.org/10.7298/h63q-9r54.

54. Software to process the microscope images and reflectance data, and to perform the GP based active learning; https://github.com/gomes-lab/SARA_ScienceAdvances.

55. G. Ashiotis, A. Deschildre, Z. Nawaz, J. P. Wright, D. Karkoulis, F. E. Picca, J. Kieffer, The fast azimuthal integration Python library: pyFAI. *J. Appl. Crystallogr.* **48**, 510–519 (2015).

56. C. E. Rasmussen, C. K. I. Williams, *Gaussian Processes for Machine Learning (Adaptive Computation and Machine Learning)* (MIT Press, 2005).

57. J. Bezanson, A. Edelman, S. Karpinski, V. B. Shah, Julia: A fresh approach to numerical computing. *SIAM Rev.* **59**, 65–98 (2017).

58. D. A. Cohn, Z. Ghahramani, M. I. Jordan, Active learning with statistical models. *J. Artif. Intell. Res.* **4**, 129–145 (1996).

59. M. Titsias, Variational learning of inducing variables in sparse Gaussian processes, in *Artificial Intelligence and Statistics* (PMLR, 2009), pp. 567–574.

60. M. Lázaro-Gredilla, M. K. Titsias, Variational heteroscedastic Gaussian process regression, in *Proceedings of the 28th International Conference on International Conference on Machine Learning ICML'11* (ACM, 2011), 841–848.

61. R. M. Neal, Monte Carlo implementation of Gaussian process models for Bayesian regression and classification. arXiv:physics/9701026 [physics.data-an] (28 January 1997).

62. Q. V. Le, A. J. Smola, S. Canu, Heteroscedastic Gaussian process regression, in *Proceedings of the 22nd International Conference on Machine Learning* (ACM, 2005), pp. 489–496.

63. K. Kersting, C. Plagemann, P. Pfaff, W. Burgard, Most likely heteroscedastic Gaussian process regression, in *Proceedings of the 24th International Conference on Machine Learning* (ACM, 2007), pp. 393–400.

64. P. Dallaire, C. Besse, B. Chaib-Draa, Learning Gaussian process models from uncertain data, in *International Conference on Neural Information Processing* (Springer, 2009), pp. 433–440.

65. E. Snelson, Z. Ghahramani, Variable noise and dimensionality reduction for sparse Gaussian processes. arXiv:1206.6873 [cs.LG] (27 June 2012).

66. A. McHutchon, C. Rasmussen, Gaussian process training with input noise, in *Advances in Neural Information Processing Systems*, J. Shawe-Taylor, R. Zemel, P. Bartlett, F. Pereira, K. Q. Weinberger, Eds. (Curran Associates Inc., 2011), vol. 24, pp. 1341–1349.

67. M. A. Álvarez, L. Rosasco, N. D. Lawrence, Kernels for vector-valued functions: A review. *Found. Trends Mach. Learn.* **4**, 195–266 (2012).

68. E. Solak, R. Murray-Smith, W. Leithead, D. Leith, C. Rasmussen, Derivative observations in Gaussian process models of dynamic systems. *Adv. Neural Inf. Proces. Syst.* **15**, 1057–1064 (2002).

69. J. Tellinghuisen, Statistical error propagation. *Chem. Eur. J.* **105**, 3917–3921 (2001).

70. B. Hu, J. Shao, M. Palta, Pseudo-r 2 in logistic regression model. *Stat. Sin.* **16**, 847–860 (2006).

71. T. J. Smith, C. M. McKenna, A comparison of logistic regression pseudo R2 indices. *MLRV* **39**, 17–26 (2013).

72. G. A. Hemmert, L. M. Schons, J. Wieseke, H. Schimmelpfennig, Log-likelihood–based Pseudo-R2in logistic Regression. *Sociol. Methods Res.* **47**, 507–531 (2018).

73. G. Gattow, D. Schütze, Über Wismutoxide. VI. Überein Wismut (III)-oxid mit höherem Sauerstoffgehalt (β-Modifikation). *Z. Anorg. Allg. Chem.* **328**, 44–68 (1964).

74. S. Hull, S. T. Norberg, M. G. Tucker, S. G. Eriksson, C. E. Mohn, S. Stølen, Neutron total scattering study of the δ and β phases of Bi_2O_3. *Dalton Trans.* **2009**, 8737–8745 (2009).

Sebastian AmentMaximilian AmslerDuncan R. SutherlandMing-Chiang ChangDan GuevarraAine B. ConnollyJohn M. GregoireMichael O. ThompsonCarla P. GomesR. Bruce van Dover

View the article online
https://www.science.org/doi/10.1126/sciadv.abg4930
Permissions
https://www.science.org/help/reprints-and-permissions

Use of think article is subject to the Terms of service