

REVIEW ARTICLE OPEN

Active learning in materials science with emphasis on adaptive sampling using uncertainties for targeted design

Turab Lookman¹, Prasanna V. Balachandran^{1,2}, Dezheng Xue^{1,3} and Ruihao Yuan^{1,3}

One of the main challenges in materials discovery is efficiently exploring the vast search space for targeted properties as approaches that rely on trial-and-error are impractical. We review how methods from the information sciences enable us to accelerate the search and discovery of new materials. In particular, active learning allows us to effectively navigate the search space iteratively to identify promising candidates for guiding experiments and computations. The approach relies on the use of uncertainties and making predictions from a surrogate model together with a utility function that prioritizes the decision making process on unexplored data. We discuss several utility functions and demonstrate their use in materials science applications, impacting both experimental and computational research. We summarize by indicating generalizations to multiple properties and multifidelity data, and identify challenges, future directions and opportunities in the emerging field of materials informatics.

npj Computational Materials (2019)5:21 ; <https://doi.org/10.1038/s41524-019-0153-8>

ACCELERATING MATERIALS DISCOVERY

A central question related to accelerating the materials discovery process is how do we guide experiments towards materials with desired properties? This is a question of optimal experimental design and our focus is to highlight how the computational and analysis tools developed in this and related areas can be employed to accelerate the materials discovery process. The Materials Genome Initiative in the U.S.,¹ where the stated objective is to reduce in half the time and costs of finding new materials, has catalyzed and spawned considerable activity. However, the number of materials discovered so far with enhanced properties, as a consequence of this initiative, is far and few between. Figure 1 illustrates a perspective in terms of our ability to find materials with increasing complexity as a function of time. Our trajectory has been one where trial-and-error has been followed by intuition, and often rapid progress is made if the synergy between theorists, who can often generate and suggest a list of compounds for possible synthesis, and experimentalists is exploited. High-throughput calculations based on $T=0\text{ K}$ first principles codes have provided the means to generate large databases of varied crystal structures and chemistries with calculated properties.²⁻⁴ Depending on the functionals used for the exchange correlations for given properties for these calculations, including stability, their fidelity will dictate how rapidly we make progress on the curve of materials complexity versus time. The thesis behind high-throughput computations is one of populating fully the phase space of materials possibilities, given the approximations inherent in these calculations, and then down selecting worthy candidates for further study.⁵ If the red curve of Fig. 1 is our desired goal (Target), the jury is very much out on whether this is an optimal strategy for the accelerated discovery process. The perspective we take in this review is that tools from statistical experimental design,⁶ pattern recognition,⁷ operations research⁸ and the fields of reinforcement and active learning⁹ in computer science offer

enormous opportunities and suggest a complementary paradigm for the discovery process. Many of these tools are unfamiliar to materials scientists but have been utilized in other fields, such as drug discovery,¹⁰ in biology for learning protein-protein interactions¹¹ and predicting macromolecular structure,¹² and cancer genomics,¹³ as well as in industry in the context of engineering design.¹⁴ For instance, in drug discovery it is not tractable to explore all possible compounds and active learning methods are used to iteratively guide experiments such that only a subset of promising experiments are carried out. The active learning methods offer, in principle, a more systematic means to approach the red line in Fig. 1 and have the potential to change the way materials science will be carried out in the future; not by changing the fundamental iterative nature of the materials design process with its reliance on distilling physical principles and mechanisms, but by improving the choices that we make to decide which material designs to test and by improving our ability to learn from previous data, experimental or computational or both. Experimental observations or calculations can be difficult, time consuming and therefore maximizing the value of experimental observation via designing experiments to be optimal by some measure is critical. We would like to be able to make decisions about what to measure, which variables to interrogate and what experimental conditions to use.

The materials challenge in its full generality encompasses a very high-dimensional discovery or search space with millions of possible compounds of which only a very small fraction have been experimentally explored. The space spans aspects of chemistry, crystal structure, processing conditions, microstructure, and the compounds can be multicomponent, for example, solid solutions and the properties can be dependent on materials descriptors or features at several length scales. Most efforts using first principles codes have largely focused on establishing a library of crystal structures and chemistries relevant to the problem, defining the

¹Theoretical Division, Los Alamos National Laboratory, Los Alamos, NM 87545, USA; ²Department of Materials Science and Engineering, University of Virginia, Charlottesville, VA 22902, USA and ³State Key Laboratory for Mechanical Behavior of Materials, Xi'an Jiaotong University, 710049 Xi'an, China
Correspondence: Turab Lookman (tll@lanl.gov)

Received: 10 April 2018 Accepted: 10 January 2019

Published online: 18 February 2019

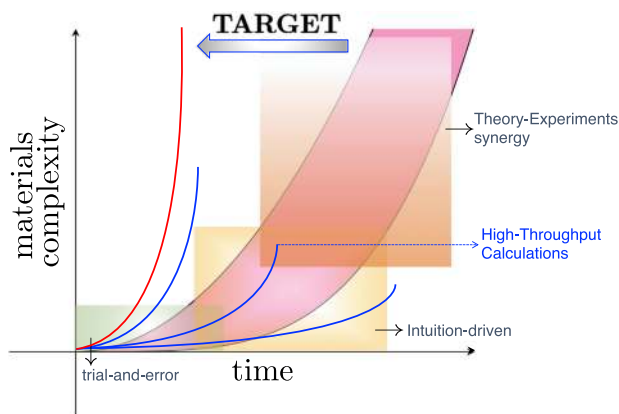


Fig. 1 The trajectory for the discovery of new materials with increasing complexity as a function of time as we accelerate the process from trial and error to using high throughput calculations and statistical design methods to improve our ability to learn from existing data and make decisions for the next materials to test (TARGET)

training space in terms of samples and features, which can be elemental properties (e.g., electronegativity, Mendeleev number etc.), bond angles, bond lengths, energetics from first principles calculations, and aspects of thermodynamics from experiments or codes such as Calphad, to down-select promising candidates for experiments or further studies.¹⁵ A number of studies have also used inference models, such as off-the-shelf regression and classification learning tools, which can include deep neural networks for large datasets of microstructures and high-throughput computational databases, to make predictions. The approach has been employed to suggest new Heusler alloys,¹⁶ polymers,¹⁷ and thermoelectrics¹⁸ to name a few. In the case of Heusler alloys, the predicted compounds have also been synthesized. In spite of a significant amount of work using this approach, few examples exist where the final properties exceed those of the best compounds in the training data sets. There are now a number of articles and reviews^{19–26} emphasizing the merits of using machine learning and statistical inference to make predictions in a large combinatorial space. However, few focus on aspects of making the optimal next decisions for synthesis and characterization by experiments or calculations.^{27–36} The predictions from machine learning are not necessarily optimal. Aspects related to multiscale modeling and constitutive response at the engineering design scale are discussed in the book by McDowell et al.³⁷

In contrast, the approaches we will discuss are based on an active learning or adaptive design paradigm, whereby the predictions from a surrogate model, which can be an inference model or a physics-based or reduced order model (ROM), serve as the input to define a utility or acquisition function, the optimal of which dictates the next experiment or calculation to be performed.^{38,39} This is the optimal experimental design component, the key decision making aspect of the active learning loop of Fig. 2. The results of the experiment or computation then augment the training data and the loop continues until the material requirements for the desired targets are met. The approach is unique in that most of the work done in the field essentially involves one or two of these steps. Taking data, building a model, making predictions and validating them with calculations or experiments. We are not aware of studies in which new materials have been found even via the inner feedback loop (green) of Fig. 2, let alone incorporating uncertainties and performing experimental design. However, a Bayesian and decision-theoretic approach discussed in this review naturally lends itself to adaptive sampling and active learning.⁴⁰ The input

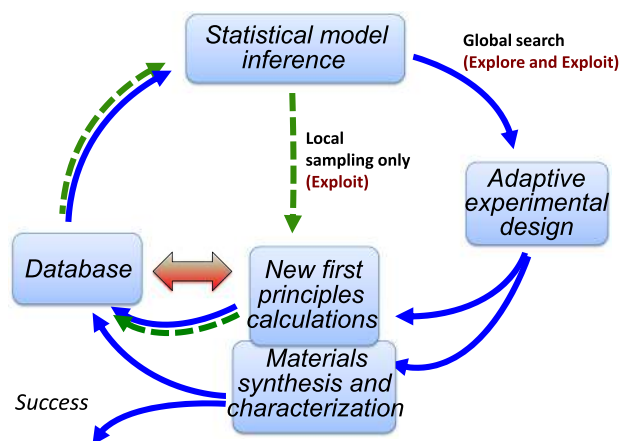


Fig. 2 The adaptive design paradigm to iteratively learn a surrogate model and use uncertainties to trade-off exploitation and exploration of the search space of unexplored materials to select the next best experiment or calculation. Red arrow represents efforts such as AFLOWLIB,² Materials Project³ and OQMD,⁴ etc

to the decision making can come from predictions from any inference, surrogate or machine learned model. One first defines the utility of an experiment or calculation and then, taking into account uncertainties in both the parameter values (if any) and the observations or objectives, chooses experiments or calculations by maximizing an expected utility. The utilities are defined according to information theoretic considerations given the desired goals. The approach can also be used for model selection, that is, the design of experiments using maximum information criteria to distinguish between models. For any experimental design procedure, we should have a notion of the value of the information (or cost of uncertainty) that is gained (or reduced) by observing a specific data point. Then, the possible alternatives to observe (experiments) can be ranked by the expected value of information they provide so that we can prioritize experiments based on this.

To set the stage, we first review briefly the classic design of experiments approach in “Design of experiments” section followed by the Bayesian approach to inference and design⁴¹ in “Adaptive sampling and Bayesian optimization” section with its focus on incorporating prior information and data to predict the property or objective. However, other approaches to inference based on machine learning methods can also be used for this purpose. The concept of utility function was previously introduced in decision theory⁴² in the context of value of information (and hence uncertainties) and this provides us with the means to make decisions about where to optimally sample next. We review a number of these functions, such as the widely studied expected improvement,³⁸ which provide criteria for making choices. This two-stage approach, also known as Bayesian Global Optimization^{40,43,44} provides a natural active learning strategy to iteratively improve the model. In “Application to materials science” section, we discuss applications of this framework to discover new compositions of materials with targeted properties, as well as increase the computational efficiency of simulations towards targeted design of optoelectronic structures, and their use in high throughput density functional theory calculations. These applications reflect areas of our own expertise, however, we follow these examples with a review of other applications and related studies. We conclude in “Challenges and future development” section by suggesting directions for future research and challenges to be overcome with its impact on materials science.

DESIGN OF EXPERIMENTS

Design of experiments (DoE)⁴⁵ encompasses many statistical approaches that seek to study the behavior of a system where there are variations in data controlling inputs and outputs. The field has a rich history dating from the work of Fisher,⁴⁶ Box and Wilson⁴⁷ and recent developments due to Taguchi.⁴⁸ In classical design, factors refer to the input variables, which can adopt many discrete levels, and their combinations over the allowed ranges of variables define the search space over which a decision needs to be made regarding the best choice of variables to use for the next decision. Factorial and Composite designs⁴⁵ are examples of the usual experimental design approach often used to find a relationship between the input and output variables. The basic idea is to distill the interrelationships between the various factors and their bearing on the response. Prioritizing these contributions and encoding them within an inference model then allows us to study how the response can be optimized. The models are usually of the form $y_i = f_{\theta}(x_i) \pm e_i$, where the uncertainties are assumed to be normally distributed, $N(0, \sigma)$. Assuming a linear model, the variance in the parameters, θ and objective, y_i , are given in terms of the moment matrix, $M = (X^T X)^{-1}$, where X denotes the design matrix with rows $(1, f_1(x_i), \dots, f_p(x_i))$ where p is the number of parameters/coefficients, such that $\text{var}(\theta) = \sigma^2 M$ and $\text{var}(y_i) = \sigma^2 M X_i$. The optimal design is then chosen based on the choice of optimality conditions, such as A - optimality which chooses x_i based on minimizing $\text{Tr}(M_x^{-1})$ and the common G—optimality, which maximizes the variance in $f(x)$.

An important consideration in DoE is to minimize the number of experiments which need to be carried out in the overall space of variables and parameters (the design space) in order to reduce costs and time. Ordinary DoE puts the same amount of resources in high and poor performance regions of the parameter space and so the costs can go up exponentially in search of the optimal design. As a result, in the early design stages where uncertainties are higher, we need appropriate strategies for the analysis of relatively expensive experiments or high-fidelity codes because seeking for an accurate numerical value of the optimum in the conceptual design stage is not efficient as there can still exist unknowns in the design that need to be addressed later in the design cycle. Also, traditional sequential design of experiments, which involves one design step followed by another, tends to be “exploitative” towards the point of interest based on previous experimental data. Hence, in the case of fairly expensive functions, adaptive sampling (also called on-line or infill sampling) is preferred over traditional design of experiments (DoE).¹⁴

ADAPTIVE SAMPLING AND BAYESIAN OPTIMIZATION

Adaptive sampling features two possible aspects that can compete: the exploration of the overall search space so that any model to be developed adequately samples globally, and exploiting local regions of the search space, where the model is likely to yield good predictions near the optimal solution. It serves the objective of rapidly assessing potential regions of interest at the initial stages of a given design. It is thus usual for engineers to employ cheap surrogate models when the engineering model is computationally intense. These surrogate models are based on a relatively small number of training points (Fig. 3) and the use of Bayesian Global Optimization (BGO) has had impact in industry for addressing large scale computational tasks since the 1990s, following the work of Jones, Schonlau and Welch, who introduced the efficient global optimization (EGO) method.³⁸ This relied substantially on early work by Kushner⁴⁹ on optimizing unknown functions, as well as work by Mockus and Zilinskas.⁵⁰ The notion of maximizing the probability of improvement from the “best-so-far” was introduced in these studies. BGO has been used extensively in design applications when optimization of unknown functions are




Fig. 3 Bayesian Global Optimization builds a surrogate model from data based on input x and selects new design points with estimates for the output \hat{y} for the next calculation, which in turn will yield the actual value y

often required, especially where the computations are expensive.¹⁴

Utility functions

The development and applications of BGO in statistics and engineering has progressed steadily and over the last five years there has been much interest from the machine learning community, where BGO has been employed for tuning hyper-parameters of machine learning models.⁵¹ Only recently have these methods been applied to problems related to materials discovery and computation, the application area we will focus on. However, complementary ideas date back to developments in decision theory with the work of Howard,⁵² who presented a theory of the value of information for the joint probabilistic and economic factors affecting decisions and how numerical values could be attributed to the reduction and suppression of any uncertainty. Lindley,⁵³ using a decision theoretic approach, brought a number of ideas together into a fully Bayesian scheme involving the definition of a utility function for experimental design so that an experiment is selected to maximize the expected utility function. The utility is defined as $U(e) = \int_Y \int_{\Theta} u(e, y, \theta) p(\theta, y|e) d\theta dy$, where $u(e, y, \theta)$ is a utility function which is chosen so that it reflects the purpose, cost or usefulness of the experiment, e , selected from a large number of possible experiments with a particular outcome, y , belonging to the set over all possible outcomes y .^{41,54} As the outcome is not known before the experiment is carried out because we do not know the parameters, θ , the expectation is taken with respect to the joint posterior distribution of θ and y . The choice of utility functions include the gain in Shannon or relative entropy (Kullback-Leibler divergence), quadratic loss and improvement from the best, which we will discuss in greater detail. We emphasize that the usually studied uncertainty quantification approaches based on variance and entropy take into account the statistical characteristics of uncertainty without regard to the objective, and the information gained associated with an experiment is measured in terms of the reduction in the entropy or variance.⁶

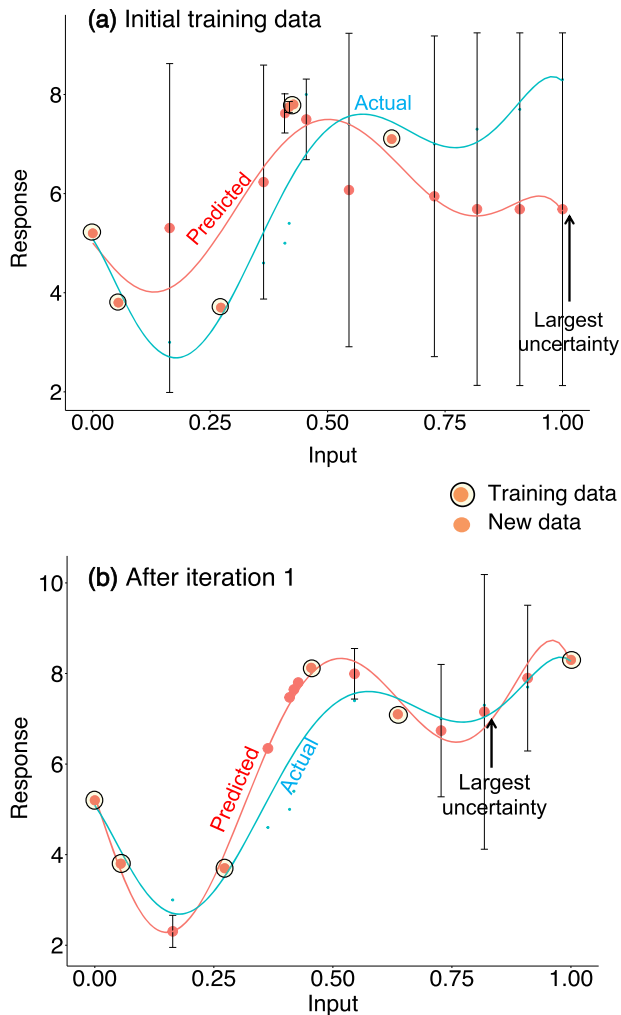


Fig. 4 Model prediction and decision making by maximizing the utility function, which in this case chooses the sample with the largest uncertainty—“exploration”. **a** Actual function (cyan) represents the ground truth function fitted to the data (Response = $-826.815x^6 + 2443.982x^5 - 2573.867x^4 + 1093.097x^3 - 118.369x^2 - 14.842x + 5.087$, where x is the Input). We randomly pick five data points (training data) to train a surrogate model with GP Regression, the predictions of which are shown by the red points (new data) and fitted curve. The GP predicts the Response (y-axis) for the test points along with the associated uncertainties (error bars). The exploration strategy recommends the data point with the largest uncertainty for the next experiment or computation. **b** The updated surrogate model after the first iteration with now 6 training data points. Notice the reduction in the error bars in **b** compared to **a**. The data point with the largest uncertainty is recommended for the next (second) iteration for validation and feedback

Two steps of Bayesian global optimization (BGO): model prediction and decision-making

The problem we want to solve is stated by $\max_x f(x)$, that is, to find the design, x that maximizes $y = f(x)$. BGO offers a powerful and efficient means for searching for extrema of objective functions which are costly and for which derivatives and convexity properties are unknown. The efficiency results from being able to incorporate prior belief in the form of smoothness, locality of $f(x)$ to guide the adaptive sampling and active learning. If $D = \{x_i, f(x_i)\}$ ($i = 1:n$) is the set of observations, then by Bayes' rule, the posterior distribution is given by $P(f|D) \propto P(D)fP(f)$, where the posterior will capture the updated beliefs about the unknown $f(x)$ given the prior information about f combined with how likelihood

the data seen so far, $P(D|f)$. This step can be interpreted as estimating $f(x)$ with a response surface or even a surrogate model, such as a Gaussian process (GP) or any machine learned model. We could choose a number of possibilities for the prior, for example, a Wiener or GP. A common example of the former is a random walk with random step sizes and this stochastic process was used in the early work of Kushner.⁴⁹ However, it has been shown that a GP, extension of the multivariate Gaussian distribution to a stochastic process with an infinite number of variables, is well suited if we make a number of simplifying assumptions and has become the mainstay for BGO.^{55,56} Instead of a Gaussian distribution over a random variable specified by a mean and variance, a GP is a distribution over functions specified by its mean function, m and covariance function, k , i.e., $f(x) \sim \mathcal{GP}(m(x), k(x, x'))$. A common choice for k is the squared exponential $k(x, x') = \exp\left(-\frac{1}{2}\|x - x'\|^2\right)$ which approaches 1 as data points x and x' get close together and 0 when they are far apart. However, the GP can allow for several kernels. In addition to the gaussian kernel $\exp(-1/2*(h/\theta)^2)$ with $h = \|x-x'\|$ and θ a parameter, the exponential $\exp(-h/\theta)$ and Matern functions $(1 + \sqrt{3} * h/\theta) * \exp(-\sqrt{3} * h/\theta)$, $(1 + \sqrt{5} * h/\theta + (1/3) * 5 * (h/\theta)^2) * \exp(-\sqrt{5} * h/\theta)$ and $\exp(-(h/\theta)^p)$ are also often used. In principle, the choice of these priors can affect the surrogate model predictions. However, this depends on the data sets. The Matern kernels are more flexible as they include more free parameters to the model compared to the standard Gaussian kernels. On the other hand, for a specific set of parameters the Matern kernel reduces to the gaussian kernel.

Figure 4 summarizes the two components of BGO:

- a model that provides the value or posterior distribution of the objective $f(x)$ at unevaluated points based on available data (x_i, y_i) with $i = 1, \dots, n$
- choice of utility function, $u(x|D_{1:t-1})$, such that its maximum over all the data in the sampling domain, $x_t = \operatorname{argmax}_x u(x|D_{1:t-1})$ provides a potentially high value of the objective at which to evaluate the function to guide the next experiment or calculation. This is the “good” decision making step to direct future sampling.

If $f(x)$ is the outcome of an experiment, the observations are influenced by noise and hence the objective needs to be written as $y(x_i) = f(x_i) + \varepsilon_i$, where the noise ε is distributed normally. The posterior distribution for $f(x)$ will now be a summation of two normal distributions and will itself be normally distributed so that the effects of noise is that the posterior distribution no longer interpolates the data points but the credible or confidence interval has positive width at all data points. In addition, the machine learning algorithms also contain “hyperparameters”, such as the penalty term, insensitive loss function term and coefficient of the kernel function, that control model complexity and help balance the bias-variance tradeoff. In general, the hyperparameters are adjusted to optimize the leave-one-out error, which in turn is computed by constructing an independent model for each training set by leaving out a single data point. The trained model now predicts the response of the left out data point. The leave-one-out error is defined as the average of the out-of-training sample residual over each model. The hyperparameters that minimize the leave-one-out error are identified and used to train a single final model. On the other hand, in a Bayesian treatment (such as GP), one can choose the hyperparameters directly from training data by minimizing the negative log marginal likelihood with respect to the hyperparameters and the noise (where the cost of Monte Carlo methods are prohibitively high). Typically, samples are drawn randomly from a prior distribution (such as the uniform prior or chosen based on domain knowledge) and

allowed to converge. The final estimate corresponds to an instance with the smallest negative log marginal likelihood.

Examples of utility functions

Expected improvement, $E[I]$. The utility function we will focus on is an improvement based function, I , initially suggested by Kushner⁴⁹ in order to maximize the probability of improvement over the best observed so far, $f(x^+)$. That is, $I = \max(f(x) - f(x^+), 0)$ would be the improvement associated with each choice x . Before evaluating $f(x)$, we do not know what the improvement will be but as we have the probability distribution on $f(x)$, we can evaluate the amount of improvement using $E[I(x)] = E[f(x) - f(x^+)]$ for positive I , where E is the expectation over the posterior distribution.³⁸ Since $f(x)$ is normally distributed with mean, μ , and standard deviation, σ , it is straightforward to show that

$$E[I(x)] = \int_{f(x^+)}^{\infty} (z - f(x^+))\phi(z)dz \quad (1)$$

$$= (\mu(x) - f(x^+)) \left[\phi\left(\frac{\mu(x) - f(x^+)}{\sigma(x)}\right) + \sigma(x)\Phi\left(\frac{\mu(x) - f(x^+)}{\sigma(x)}\right) \right]$$

where ϕ and Φ are the standard normal density and cumulative distribution functions. Thus, a new sample point, x^* , is chosen amongst other data points based on the largest expected improvement, i.e. $x = \operatorname{argmax}_x E[I(x)]$. The limiting cases of uncertainty on $E[I(x)]$ give us:

- Small σ : $E[I(x)] \rightarrow \mu - f(x^+)$, i.e. choose the μ greater than $f(x^+)$ to maximize $E[I(x)]$ or exploit the model predictions.
- Large σ : $E[I(x)] \rightarrow \sigma$, i.e. choose the μ with largest uncertainty to maximize $E[I(x)]$ or explore the sample space.

Additionally, $E[I(x)]$ is largest where both $\mu(x)$ and $\sigma(x)$ are large. That is, points likely to give the most gains are those with the best prediction but also where uncertainty is greatest. These aspects emphasize the *exploration—exploitation* tradeoff, which is central to aspects of active⁵⁷ and reinforcement learning.⁵⁸ It characterizes many problems, including the multi-armed bandit problem, in which decisions need to be made repeatedly but where the outcomes are uncertain and we wish to obtain immediate results as to where to sample optimally so that we get better results in the future. The multi-armed bandit problem⁵⁹ is an example illustrating the exploration-exploitation dilemma and refers to a slot machine with n -arms or “bandits” with each arm having its own probability distribution of success. Pulling any one of the arms gives a stochastic reward of either success or failure. The objective is to pull the arms one-by-one in sequence such that the total reward collected in the long run is maximized. The problem has applications in many real-world situations where one would like to select the “best” bandit out of a group of bandits i.e., A/B testing, line-up optimization, evaluating social media influence.

Knowledge gradient (KG). In the presence of noise in the data, the utility function becomes a generalization of expected improvement, namely knowledge gradient, $KG(x)$.^{43,60} With noise, the $f(x)$ values are not exactly known. Instead, we consider $\mu(x)$ so that best choice at the next step with the largest improvement is given by $\mu_{n+1}^* = \max_x \mu_{n+1}(x)$ and $KG(x) = E[\mu_{n+1}^* - \mu_n^*]$. If we normalize the change in means by the number of standard deviations, σ , then $KG(x)$ is given by,

$$KG(x) = \sigma \left[\phi\left(\frac{\mu(x) - f(x^+)}{\sigma(x)}\right) \right] + \Phi\left(\frac{\mu(x) - f(x^+)}{\sigma(x)}\right) \quad (2)$$

The next sampling point is then one for which $KG(x)$ is largest, that is $\operatorname{argmax}_x KG(x)$. The sampling region over which μ_n^* , μ_{n+1}^* are determined is not restricted to the observed data, as in the case of $E[I(x)]$, but can be extensive covering many allowed possibilities. This aspect, and the noise in the data, are the key difference from

expected improvement. Recently, expected improvement has been generalized to also include noise.⁶¹

Mean objective cost of uncertainty (MOCU). An objective based uncertainty quantification scheme was recently introduced to study interactions among genes to develop drugs to mitigate aberrant phenotypes, such as cancers, in the design of gene regulatory networks.^{62–64} The basic idea is that the presence of large uncertainties can degrade experimental design and the aim of new measurements should be to reduce uncertainties in finding materials with desirable properties. That is, one needs an objective-based uncertainty quantification scheme as compared to reducing overall variances or entropy to monitor information gain in an experiment. A prior distribution that encodes knowledge about the unknown parameters or features, say θ , is assumed, which after knowing the outcome of an experiment, is updated to a posterior distribution. If further measurements are required, then this distribution is used as the new prior for the next experimental design loop. The MOCU serves as the utility function and measures the deterioration in the design due to the presence of model uncertainty, i.e., it measures the degradation in performance between an optimal design based on partial prior knowledge and data compared to a design with full knowledge of the system. One chooses the next experiment to shrink the uncertainty the most, i.e., the experiment that is expected to minimize the variance in the posterior distribution the most. That is, the alternative that is expected to decrease the MOCU the most is selected at each iteration, so that after observing that experiment’s output, we would be able to make a better decision towards the objective.

Clarifying further, suppose we are at iteration 0 and we have a cost function $f(x)$ with unknown parameters θ , i.e., we have observed initial data points and have updated prior distributions to posterior, $f(x)$, but not yet selected the next experiment x . At this step if we want to stop doing any more measurements and select the x that, for example, maximizes $f(x)$ based on our current state of belief, we would select x which maximizes the expected value of $f(x)$ over the unknown parameters, θ , assuming we have a prior distribution for θ to allow us to evaluate the expected value. Thus, we would select $x_{\text{robust}} = \operatorname{argmax}_x E_{\theta} f(x)$. This is the best we can do “on average”, hence the term “robust”, because we do not know the θ ’s and we are not guaranteed to choose the true optimal. In many ways it is this “robust” we seek which is expected to perform optimally *on average*. Now, corresponding to each parameter value, θ , the “optimal” material, x^+ , can be found by performing $x^+ = \operatorname{argmax}_x f_{\theta}(x)$. We will have an optimal for each parameter value θ and this is our best choice if there are no unknown parameters. As we do not know the true parameter values, θ , the amount we lose is the difference between the true maximum value $f(x^+)$ and the value associated with the robust selection. This difference in the absence and presence of uncertainty is the objective cost of uncertainty (OCU) and its expectation over all θ is MOCU, which is given by

$$\text{MOCU} = E_{\theta}[f_{\theta}(x^+) - f_{\theta}(x_{\text{robust}})] \quad (3)$$

The next experiment selected is that which reduces MOCU the most. Hence, for each x , we consider all possible values of the experimental outcome, y , to calculate MOCU weighted by the probability distribution, $P(y|x)$, of y being the true experiment outcome of x . The selected experiment for measurement, x^* , is then given by

$$x^* = \operatorname{argmin}_x E_{y|x}(E_{\theta|x}[f_{\theta|x}(x^+) - f_{\theta|x}(x_{\text{robust}})]), \quad (4)$$

where $E_{y|x}$ is expectation over $P(y|x)$. If we want to stop measurements at the next time step and select the candidate with highest expected $y = f(x)$, our final selection will be given by the robust alternative by taking the expectation over θ .

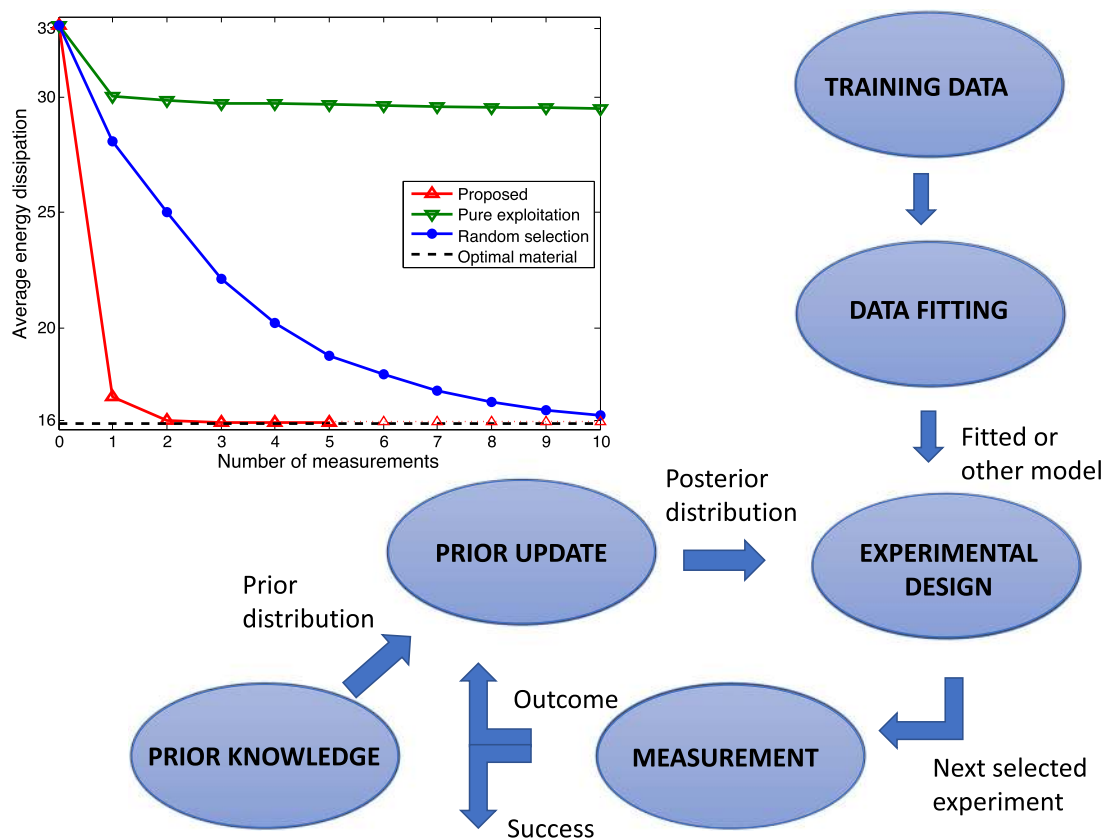


Fig. 5 The objective based uncertainty approach, MOCU, to experimental design (right) with its emphasis on construction of a prior distribution to commence the search. The plot on the left shows the performance of the method for simulations of shape memory alloys where the objective, the average energy dissipation, is plotted versus the number of measurements.²⁸ MOCU finds the material features for minimizing the objective in far fewer experiments than pure exploitation or random selection. Reproduced with permission from ref. ²⁸ Copyright Elsevier 2017

This method has recently been applied within Landau or phase field simulations to determine dopants that minimize the energy dissipation, which affects the fatigue of shape memory alloy such as FePd.²⁸ These simulations assume a random initial dopant distribution in the FePd matrix and associated with the dopants is a concentration, potency and range, all of which can vary. Figure 5 summarizes the active learning strategy that the authors have used leveraging MOCU. In the simulations, a stress is incrementally applied to obtain a stress versus strain curve from which the energy dissipated is estimated. The data from the simulations serve as the training data to essentially fit a nonlinear model to the dissipation in terms of the concentration, potency and range of the dopants. This model is then used for the experimental design step in Fig. 5 to select the next best dopant and its features on the basis of MOCU, that is, the dopant that minimizes the overall deterioration in the model uncertainty. After the measurement is performed (in the real situation) or as here, after the outcome is known, the prior distribution in terms of the unknown dopant parameters is then updated to give a posterior distribution, which serves as the new prior for the next step in the active learning loop.

The prior distribution is chosen as a Dirichlet distribution by generating different samples using different dopant concentrations, strength and parameters. The plot in Fig. 5 shows how the average energy dissipation behaves after different numbers of measurements. The five measurements for the proposed approach are compared to ten measurements for pure exploitation and random selection selection policies. A summary of the utility functions discussed here is given in Table 1

APPLICATIONS TO MATERIALS SCIENCE

We discuss three independent examples to demonstrate the breadth of application of the active learning strategies in materials science. In all three examples discussed below, several iterations of the design loop are performed with five iterations for each of the *four* AL strategies of the experimental materials discovery of ferroelectrics in Example 1, up to a 1000 iterations for the optoelectronic simulations of Example 2, and twenty five iterations for *one* AL strategy of the first principles code used in Example 3. Unlike the more common statistical studies (discussed later) on data with assumed distributions, experimental data as well as computational data from physics codes do not offer the luxury to randomly select different training sets because each evaluation of a sample data point is expensive.

Materials discovery

The discovery of materials with targeted properties typically involves investigating a vast search space efficiently. Systematic design strategies aid in guiding or recommending iteratively the optimal compounds for synthesis. The strategies discussed above have been applied to alloys^{27,29} and ceramic materials.^{30,35,36}

Example 1: piezoelectrics with large electrostrains. We review a recent example³⁵ where the goal was to find Barium Titanate (BaTiO₃) based piezoelectrics with large electrostrains. The idea is to use chemical substitution (i.e., doping), such that Ca²⁺ and Sr²⁺ cations replace Ba²⁺ on the A-site of the ABO₃ perovskite structure, and Zr⁴⁺ and Sn⁴⁺ substitute for Ti⁴⁺ on the B-site.

| Utility function | Description |
|------------------------------------|---|
| Efficient global optimization | Evaluates trade-off between exploration and exploitation in the <i>absence</i> of uncertainties in the response |
| Knowledge gradient | Evaluates trade-off between exploration and exploitation in the <i>presence</i> of uncertainties in the response |
| Mean objective cost of uncertainty | Evaluates the degradation in performance between an optimal design based on partial prior knowledge and data compared to a design with full knowledge of the system |

A summary of the three main utility functions discussed in this review

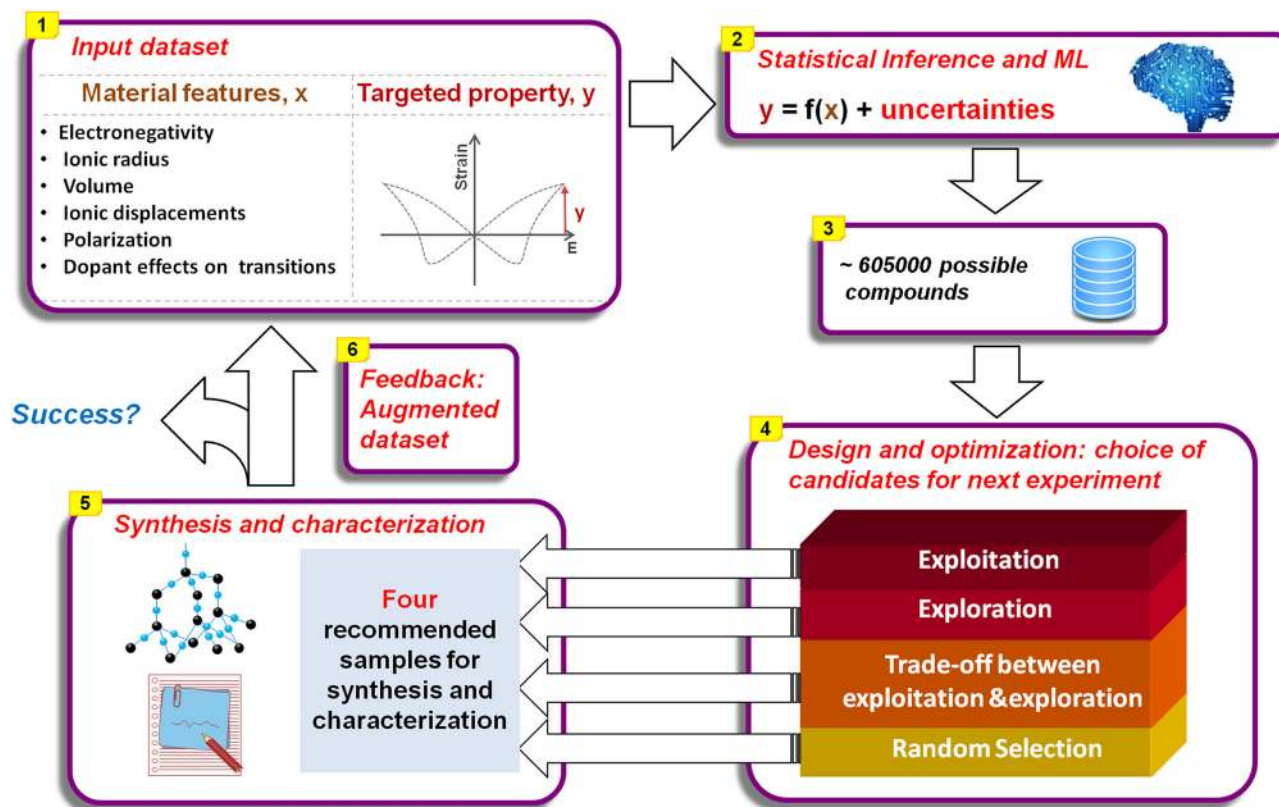


Fig. 6 Active learning loop comparing experimentally the performance of the expected improvement, $E[I]$ (trade-off), based on balancing the trade-off between exploitation-exploration with design selection strategies based on **a** maximum uncertainty from the surrogate model prediction **b** the best (maximum) prediction from the model, and **c** purely random choice, for finding a BaTiO_3 based piezoelectric with the largest electrostrain at 20 kV/cm. Reproduced with permission from ref. ³⁵ Copyright John Wiley and Sons, Inc. 2018

Thus, the family of solid solutions being considered is given by $(\text{Ba}_{1-x-y}\text{Ca}_x\text{Sr}_y)(\text{Ti}_{1-u-v}\text{Zr}_u\text{Sn}_v)\text{O}_3$, where the compositions x , y , u , and v are constrained by $1-x-y > 0.6$, $x < 0.4$, $y < 0.3$, $1-u-v > 0.6$, $u < 0.3$ and $v < 0.3$ so that relaxor phases are avoided. As the composition of each ion can be controlled to 0.01 in the synthesis process, there are potentially about 605,000 possible compositions of which the training data consists of only 61 which have been synthesized (~0.01%). All 61 training data points, corresponding to different chemistries and compositions, were synthesized and characterized using the identical protocols by the same researchers to remove lab-to-lab variability and processing. In addition, all data points were chosen essentially at random, i.e., no prior knowledge is used in performing the experiments or correlating the results from differing compositions to build the training data set.

The features, which should be easily available for all unexplored materials, are expected to be physically meaningful and should have a bearing on the property. The seven features (Fig. 6) we employed included microscopic attributes such as electronegativity, polarizability, ionic displacement, ionic radius and volume,

as they impact polarization and strain, as well as features which capture the direction (increase, decrease or no change) of the dependence of the cubic to tetragonal ferroelectric transition temperature and tetragonal to orthorhombic ferroelectric transition temperature, respectively, on the doping elements. The data has been uploaded as part of the Supplemental Information.

The result of the dimensional reduction using multidimensional scaling is given in the Supplementary Fig. 1, which shows little evidence of clustering amongst the compounds. Moreover, the analysis of the way the ML was done, including feature selection, the use of bootstrap methods for ensemble sampling to estimate uncertainties of the surrogate model, are given in ref. ³⁵ and the Supplement to it, which discuss these aspects. A number of models, including polynomial fits, Lasso support vector machine (SVR) with a linear (SVR.LIN) and radial-based (SVR.RBF) kernel function, and Gradient Tree Boosting were used. Given such a high quality experimental data set from an unknown distribution, the statistical method of non-parametric bootstrap sampling with random replacement⁶⁵ was used to generate 1000 samples of training data. The basis of bootstrap sampling is that the empirical

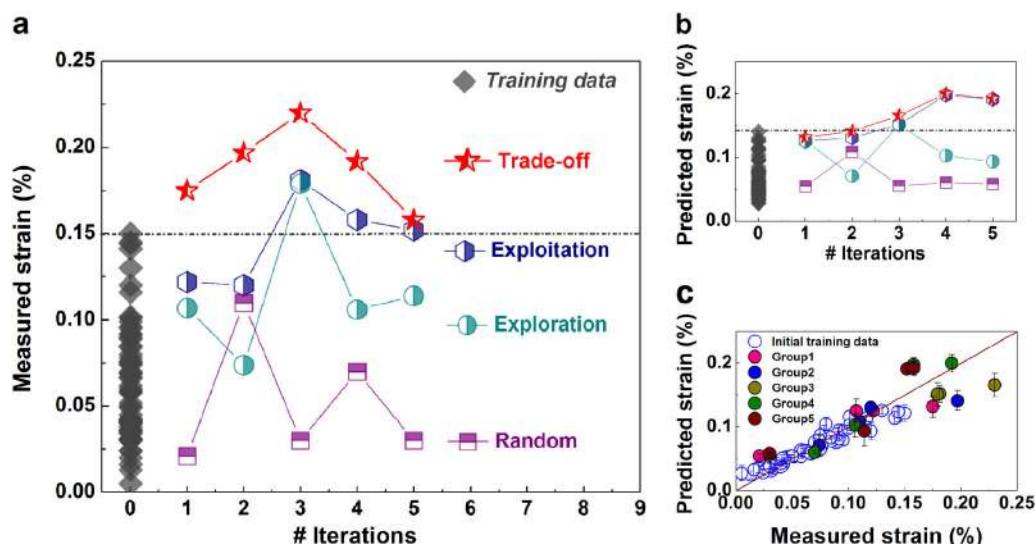


Fig. 7 **a** Comparison of the relative performance of the design on the four selection strategies after synthesis and characterization of new compounds after five iterations of the loop. The third iteration leads to the compound $(\text{Ba}_{0.84}\text{Ca}_{0.16})(\text{Ti}_{0.90}\text{Zr}_{0.07}\text{Sn}_{0.03})\text{O}_3$ with the largest electrostrain with $E[\text{I}]$ the best performer in every iteration. **b** The predictions of the design. **c** The quality of the model based on the training data and after each iteration. Reproduced with permission from ref.³⁵ Copyright John Wiley and Sons, Inc. 2018

distribution so obtained mimics the statistics associated with the original population. An ensemble of 1000 models then gave predictions with means and variances. Cross-validation (CV), typically 10-fold and grid-search to optimize the hyperparameters, was used for each model. To be specific, for the case of SVR on the given training data, we first constructed a model with the best combination of hyperparameters leading to the smallest CV error. We then evaluated the model performance by monitoring two errors, the cross-validation error and the mean squared error. For the mean squared error, we used bootstrap sampling on the whole training data with replacement. For the selected bootstrap sample, one model was built with the hyperparameters tuned by 10-fold cross-validation, and then applied on the whole training data. In total we had 1000 bootstrap samples, so that 1000 models were trained giving 1000 predictions, from which the mean squared error was calculated.

A number of steps were carried out to build an adequate model. Before the first iteration, a model was built with 55 data points and 6 randomly chosen test points using bootstrap sampling and 10-fold cross validation on the 55 points. These results together with the test results on the 6 points are shown in Supplemental Fig. 2. The MSEerror for the training data was 0.0047 and that for the test data was 0.0044. This suggests that the model is not overfitting. Also considered were safeguards usually undertaken to avoid overfitting, including regularization as well as the use of ensemble models. Regarding the former, the hyperparameters were optimized using cross validation and for the latter, bagging as well as boosting were used (see Supplemental Fig. 3 for the model fit on the 55 points using SVR.RBF with bootstrap and Gradient tree boosting), and they give similar error estimates.

Figure 6 captures the active learning loop in which the performance of the expected improvement, $E[\text{I}]$ is compared with other design selection strategies experimentally. These include selection of a composition for synthesis from all allowed possible samples based on (a) maximum uncertainty in the electrostrain from the surrogate model prediction (b) the best (maximum) prediction from the model (c) purely random choice. Thus (a) corresponds to pure exploration in Fig. 6, (b) to pure exploitation, and the trade-off between the two is characterized by $E[\text{I}]$. Random bootstrap sampling (ref.⁶⁵) with replacement on the training data set was used to generate 1000 samples, from which 1000 models were learned to predict the mean and standard

deviation. Figure 7a compares the relative performance of the design after synthesis and characterization of the new compounds predicted by each of the selection criteria after five iterations of the loop. After each iteration, the new results augment the training data and the iterative loop repeats. The compound with the largest electrostrain, $(\text{Ba}_{0.84}\text{Ca}_{0.16})(\text{Ti}_{0.90}\text{Zr}_{0.07}\text{Sn}_{0.03})\text{O}_3$, was obtained on the third iteration. The trade-off, given by $E[\text{I}]$, was the best performing strategy in every iteration with random selection being the worst overall performer. Interestingly, pure exploration and exploitation also had their peaks in the third iteration. The predictions in Fig. 7b of the electrostrain from the inference and different selection criteria show a similar tendency to the experimental results, although the compound with the largest strain occurs an iteration later.

Figure 7c shows the measured strain (x-axis) vs predicted strain (y-axis) obtained by SVR.RBF. Groups 1 to 5 indicate the 20 new compositions that were added as a result of AL (four per Group or iteration). The blue data points correspond to the original trained model before any new experiments. The model at the end of the fourth group of experiments is shown in Supplemental Information 4; the experimental measurements, overall, are quite consistent with the predictions. Pure exploitation is superior to pure exploration. It tends to sample the property-feature space in the vicinity of local minima, in contrast to exploration which samples more globally in regions where there is limited data and uncertainties are high, and is preferred for this problem. As 9 of the 20 compounds had a strain value larger than the best in the training set, the Fisher p value <0.001 , indicating that the result is significant and the probability that this result is due to chance, is low. Also, the results of experiments (as depicted in Fig. 7a) show that the AL strategies perform better over five experimental iterations and 20 newly synthesized/characterized samples than random sampling. Additionally, Supplementary Figure 5 shows that for the original training data set, a bootstrap statistical analysis with replacement yields that the number of new measurements required to find the sample with the largest electrostrain as a function of training size, is far fewer for AL strategies than random sampling. This result is consistent with the results of Fig. 7a from *in-fact* experiments.

We observe from Fig. 7c (and also from Fig. 12b of Example 3) that although the initial surrogate model is reasonable, subsequent iterations are not particularly improving the model, if

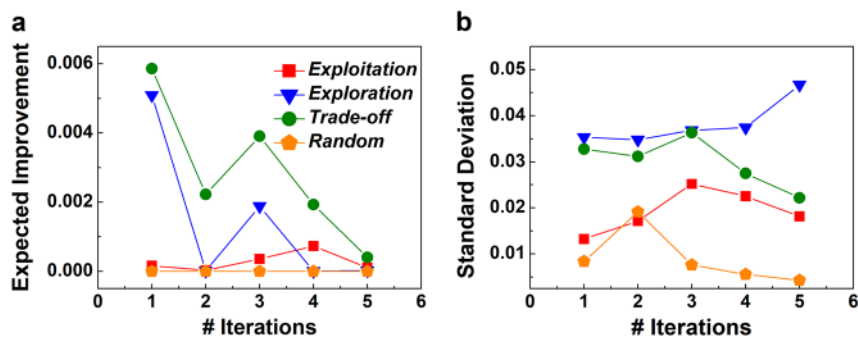


Fig. 8 The behavior of $E[I]$, the expected improvement, and the standard deviation as a function of the number of experimental loops carried out. **a** The expected improvement does not monotonically decrease, but the trend is a decreasing as the expected improvement will eventually decrease from further experiments. For the case of the randomly selected compounds, their $E[I]$ is essentially zero as most of the compounds in the search space of possibilities have an $E[I]$ that is very small or close to zero. **b** The uncertainty or standard deviation in the surrogate model decreases with iteration number, except for pure exploration, which is driven by selecting compounds with maximum variance

anything the model can be argued to get worse as the design guides the selection of the next compound for synthesis. More work is needed to capture the underlying reasons and to study the interplay between the quality of the surrogate model and the observed behavior.

Figure 8a, b show the performance of the design as a function of the number of iterations of the loop. In Fig. 8a, we plot the calculated expected improvement (y-axis) for each composition chosen from each AL strategy as a function of five iterations (x-axis). The expected improvement does not monotonically decrease; it fluctuates but the trend is a decreasing one so that eventually the expected improvement will decrease from further experiments, and in the limit we expect it to tend to zero (a more rigorous study with computational data is discussed in Section IV B. 2). In Fig. 8b, we also plot the corresponding standard deviation (our measure of prediction uncertainty) as a function of the five iterations of experiments. The standard deviation of the selected samples, together with their means, goes in determining $E[I]$. For $E[I]$, the standard deviation decreases with increasing iterations as exploitation of the model is more favored. For the case of the randomly selected compounds, their $E[I]$ values are essentially flat. This is because most of the compounds in the $\sim 605,000$ possibilities have an $E[I]$ that is very small or close to zero (that is, most compounds are not likely to lead to any improvement). As expected, the standard deviations are relatively large for the compounds selected on the basis of the pure exploration strategy.

Simulations for targeted design

We focus here on two examples that discuss the application of active learning methods to iteratively guide physics-based computational codes, such as a Poisson-Schrödinger equation solver (APSYS) used for optoelectronic simulations to design structures of light emitting diodes (LEDs), and density functional theory (DFT). The objective is to build cheap, surrogate models that are sufficiently accurate in carrying out specific tasks without the need for executing these expensive physics-based codes every time a new instance is chosen. We discuss the details below.

Example 2: building surrogate models for optoelectronic simulations.

Poisson-Schrödinger (P-S) simulations are often used to optimize fabrication of light emitting diode (LED) structures. The structures are built from a series of quantum wells and the objective is to maximize the optical output, typically the quantum efficiency or yield.⁶⁶ Such simulations for optoelectronic applications self-consistently solve coupled classical and quantum equations and are available in commercial packages such as APSYS. The process is quite time consuming as different inputs to the code are tried

until a satisfactory structure or conditions for fabrication are obtained. The simulations themselves are computationally demanding and thus what is desired is to rapidly construct a surrogate model that maps the LED structure or inputs to simulated efficiency, thereby overcoming time-consuming trial and error based simulations to obtain the optimal design in as few iterations as possible. Leduc et al.³¹ addressed this problem by starting with a database of LEDs with known structure (i.e., number of layers, their composition, doping, and widths) with labeled electro-luminescence internal quantum efficiency, and built a GP surrogate model to make predictions of the efficiency for as-yet unseen structures. They used the $E[I]$ algorithm for the selection and ranking and their work shows that to find a globally optimum LED, striking a delicate balance between exploration and exploitation is favored in selecting new sample points. This trade-off serves to improve the global accuracy of the model, thereby minimizing the possibility of becoming stuck in a local region of the LED design space for which the efficiencies are only locally optimal.

Figure 9a shows the structure of the GaN LED studied by Leduc et al. consisting of a feature space of 6 dimensions, so that each point, x was defined by the Indium composition of each of the 5 quantum wells as well as the average Indium value of the barriers. The width of the quantum wells will be a function of the indium levels as well as the barrier height so that the wavelength does not especially vary. The input parameters used in the APSYS code were those from ref.⁶⁶ with a current density of 75 A/cm^2 and the output included the band structure as well as the results of the solution of the coupled transport equations subject to the band structure deformation and carrier redistribution. Figure 9b shows the results for the internal quantum efficiency calculated using APSYS with GP and $E[I]$ performing the optimization. Figure 9b (i) shows that within 25 iterations, the code has reached 75% of the total efficiency, hence BGO converges rapidly, and after 75 iterations the efficiency is at its optimal with little improvement thereafter (Fig. 9b (ii)). From 150 iterations onwards, the simulation is essentially working towards reducing the uncertainties so that at the end, R^2 , the coefficient of determination, as determined by cross validation, is greater than 0.99 (Fig. 9b (iii), (iv)).

Poisson-Schrödinger simulations require tuning of a number of parameters, such as the damping step of the Newton solver that reaches equilibrium, and convergence is often an issue. Optimizing parameters is a general problem in computational sciences and physics^{55,67–70} and in machine learning.^{71,72} Leduc et al.³² recently demonstrated how a surrogate model, using random forest (RF) to predict the convergence rates and $E[I]$ to do the sampling, can be applied to speed up convergence of simulations, such as the P-S equations encoded in APSYS.

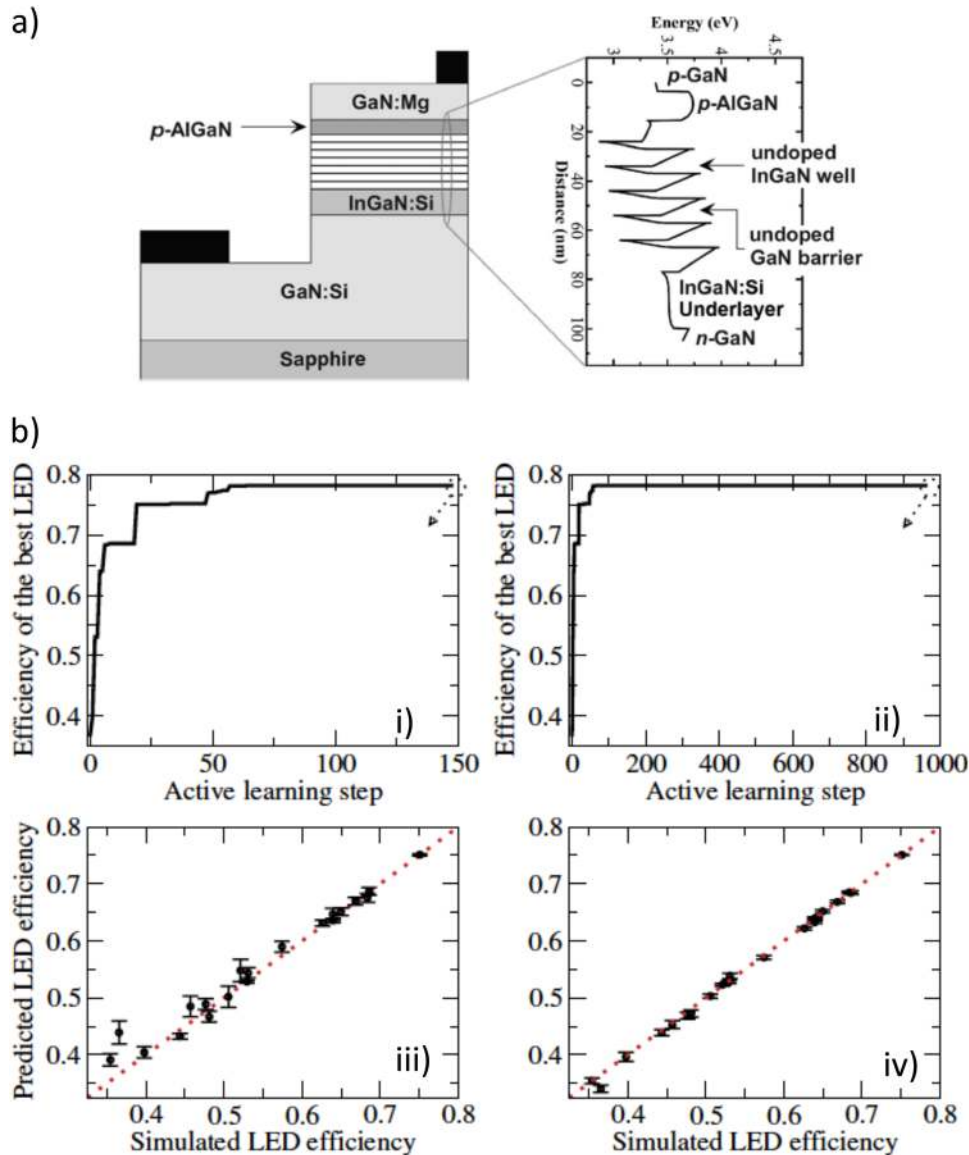


Fig. 9 **a** A schematic of the structure of a light emitting diode (LED) showing quantum wells and conduction band that was used in the surrogate based modeling and optimization **b** (i) and (ii) Results of the simulated quantum efficiency of the LED as a function of the active learning steps executed by the code. The simulation converges rapidly and there is little improvement beyond the optimal reached, except that the uncertainties decrease. (iii), (iv) Parity plots of the out of sample data showing the efficiencies predicted by the GP model and those obtained from the simulations after (iii) 150 iterations and (iv) 1000 iterations when the uncertainties are much smaller and the fits superior. Reproduced from ref. ³¹

Rather than a GP model, Leduc et al.³² utilized a RF model to predict the fraction of convergence for a given simulation, so that 0% refers to a simulation that does not converge at all and 100% identifies a simulation successfully converged. For this application, the completion fraction of optoelectronic simulations in which voltages or currents are ramped up to target values, determines the convergence rate. For example, a simulation where the drain to source voltage is ramped up from 0 to 10 V, but the simulation fails for 7.5 V, is said to have converged at 75%. A RF is a collection of decision trees, and the prediction of the forest is an average of the predictions of each tree, hence various measures of uncertainty can be evaluated. In this work, the empirical estimate of the error between the RF prediction and the true convergence ratio was used as an uncertainty measure.

Figure 10 demonstrates the efficacy of the approach to parameterization. The first 15 iterations show the lack of

convergence of simulations for which parameters have been chosen randomly. The simulations that follow from the black line onwards in Fig. 10, using a database of the initial random simulations, are then set up automatically using the active learning strategy with RF and $E[\cdot]$. With the points showing the convergence fraction, these simulations quickly converge, and similar to the simulations in Fig. 9b, there is little subsequent improvement except the uncertainties decrease. The occasional drops in the convergence fraction that are seen in Fig. 10 are indicative of the exploration of regions in parameter space as the simulation proceeds toward building a robust model. The coefficient of determination, R^2 , past 150 iterations is ~ 0.9 . This example on GaN simulations demonstrates the potential of machine learning for automated speed up in search of target properties as well as convergence.

Example 3: guiding density functional theory codes. One of the popular approaches in computational materials science is to use density functional theory (DFT) methods to calculate properties such as the band gap (E_g), thermal conductivity, and modulus (elastic, bulk, and shear) to name a few, in a high-throughput manner.^{2–4,17} More recently, there is also interest in exploring the potential use of machine learning methods to learn from DFT data and, in turn, make new predictions.^{17,73–76} Typically, machine learning combined with DFT studies has the following flavor: a reasonably large dataset from high-throughput DFT is built, a machine learning model is trained on the dataset, and the trained model then predicts new responses for compositions not present in the dataset. Note that this approach does not account for

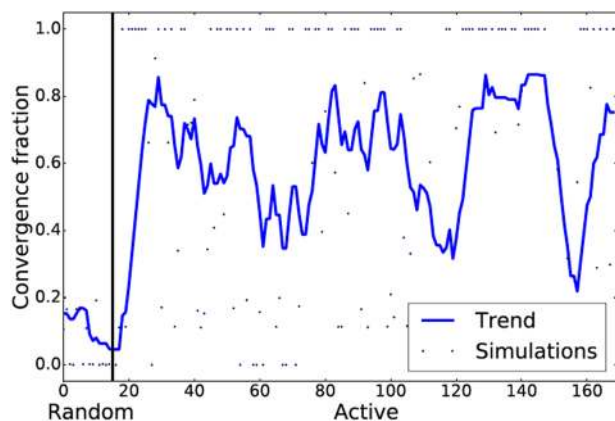


Fig. 10 Active learning applied to the convergence of GaN transistor simulations by optimizing numerical parameters (5 in total). Each point on the plot is the convergence ratio for a simulation, and the blue line shows the moving average or trend of the convergence ratio. A random choice of numerical parameters leads to an average convergence fraction of 10%, with no simulations fully converging. These random simulations serve as an initial database and as the active learning is initiated (black vertical line), the simulations converge fully within a few iterations. Reproduced with permission from ref. ³² Copyright AIP Publishing 2017

uncertainties or active learning. Here, we demonstrate a specific application of active learning where machine learning methods iteratively guide DFT calculations towards promising regions in the composition space.⁷⁷

We consider stoichiometric apatite compounds for this work that have a chemical formula of $A_{10}(BO_4)_6X_2$, where A and B are divalent and pentavalent cations, respectively, and X is an anion. Apatites find applications as biomaterials, luminescent materials, and host lattices for immobilizing heavy and toxic elements and radiation tolerant materials.⁷⁸ We focus on a chemical space, where $A = \{\text{Mg, Ca, Sr, Ba, Zn, Cd, Hg, or Pb}\}$, $B = \{\text{P, As, or V}\}$, and $X = \{\text{F, Cl, Br, or OH}\}$. This gives a total of 96 unique compositions that satisfy the $A_{10}(BO_4)_6X_2$ stoichiometry. The goal is to rapidly find an apatite composition with the largest E_g by combining active learning methods with DFT calculations in as few iterations as possible.

The training data consists of 13 randomly chosen apatites for which E_g is calculated within the generalized gradient approximation (GGA). Each apatite is represented by r_A , r_B , and r_X , which correspond to the Shannon ionic radius of A-site, B-site, and X-site elements, respectively.⁷⁹ The largest E_g (5.35 eV) in the training set belongs to the $\text{Sr}_{10}(\text{PO}_4)_6\text{F}_2$ (SrPF) composition in the $P6_3/m$ crystal symmetry. We used an ensemble of 100 SVR-RBF ML models to establish a relationship between the three features (r_A , r_B , and r_X) and the property, E_g . The hyperparameters for each SVR-RBF model were optimized using the leave-one-out cross-validation method. Each SVR-RBF model will return a prediction for E_g . Since there are 100 such models, we can estimate the mean (μ or \hat{E}_g) and standard deviation (error bar, σ) from the 100 E_g predictions. We estimated the resubstitution absolute mean squared error to be 0.54 eV/composition.

The iterative design loop is shown in Fig. 11a. Our training set contained 13 compositions from which we built an ensemble of 100 SVR-RBF models. Our next step is to predict the $\hat{E}_g \pm \sigma$ for the remaining 83 compositions for which the E_g data from DFT-GGA calculations is not known. We recommend a composition that has the largest $E[I]$ for DFT-GGA validation and feedback. The training set is then augmented with this new composition. We set a computational “budget” of 25 total iterations to gain an understanding of the adaptive design process. In Fig. 11b, we show the DFT-GGA calculated E_g data for the compositions recommended.

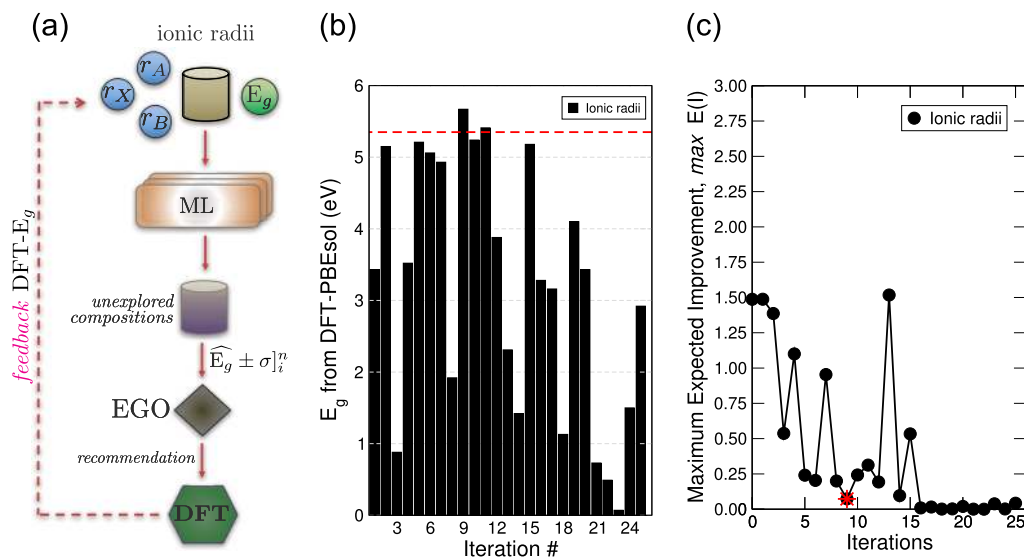


Fig. 11 Active learning for guiding Density Functional Theory (DFT) calculations. **a** Overarching strategy in iteratively guiding DFT calculations to find a composition with the largest band gap. **b** DFT-GGA E_g data (y-axis) as a function of iteration number (x-axis), where the red dashed line represents a composition with the largest E_g in our initial training set. **c** Evolution of maximum $E[I]$ at the end of each iteration showing that it does not decrease monotonically, but fluctuates. Further, the data point corresponding to the best composition (9th iteration) need not coincide with the largest $E[I]$ in the whole composition space. Adapted with permission from ref. ⁷⁷ Copyright Springer 2018

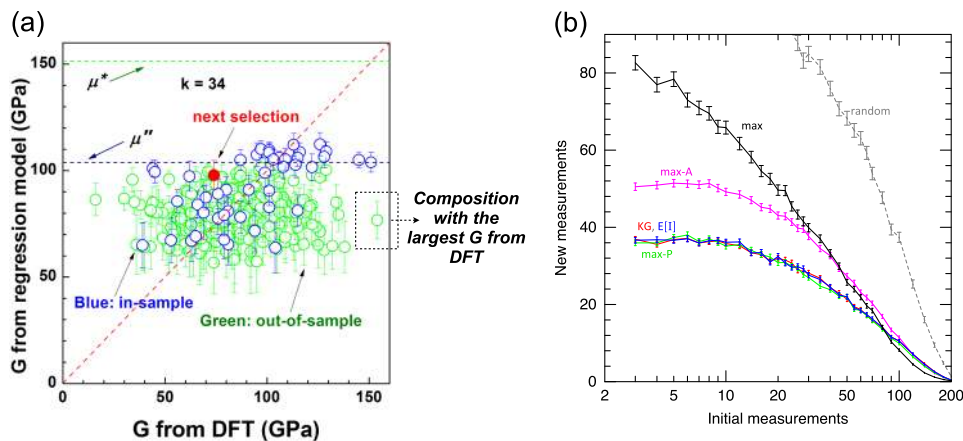


Fig. 12 Model quality and active learning. **a** Performance of the SVR model (y-axis) vs. DFT data (x-axis) in predicting the shear modulus (G) of MAX phases. The blue and green data points correspond to in-sample and out-of-sample data points. The red data point correspond to the next recommendation from active learning at the “ k ” = 34th measurement. Symbols μ^* and μ'' represent data points with the largest G in the population and SVR predictions on the out-of-sample data points. **b** Relative performance of various active learning methods (including random sampling, max: highest expected score, max-A: alternates between choosing the highest expected score and the most uncertain estimated score, max-P: maximizes the probability of improvement) in maximizing G . The smaller the number of new measurements (y-axis), the better the accelerated search. X-axis shows the number of initial in-sample data points used for training the SVR model. Error bar is the standard deviation from 1000 random trials. Reproduced from ref. ⁹⁹

The optimal composition, $[\text{Ca}_{10}(\text{PO}_4)_6\text{F}_2]$, with the largest E_g of 5.67 eV was identified in the 9th iteration.

In order to provide insights into the iterative adaptive learning process, we also tracked the $\text{max}E[I]$ at the end of each iteration. Ideally, we expect $\text{max}E[I]$ to have large values during the initial iterations and to monotonically decrease as the number of iterations increases (especially after the optimal material is identified). However, we find that $\text{max}E[I]$ shows a non-monotonic trend and does not decrease smoothly as a function of the number of iterations. This is shown in Fig. 11c. The $\text{max}E[I]$ consistently attains a value of zero only from the 16th iteration onwards. This is intriguing because the optimal composition was found in the 9th iteration, yet the $\text{max}E[I]$ did not reduce to zero. This outcome sheds light into the one of the difficult questions related to the stopping criterion, i.e., when should one stop the iterative loop? We do not recommend stopping immediately after $\text{max}E[I]$ has reached a small value. Instead, it is important to confirm that the $\text{max}E[I]$ is consistently small and does not increase. This can be accomplished by running additional experiments or simulations. An alternative $\text{max}E[I]$ -agnostic criterion would be to stop the iterative loop when a material with the desired response is found. However, this approach will not ensure an optimal ML model because the search space will not be surveyed adequately.

Review of other applications and studies

In addition to examples we have already referred to, there have been a number of recent efforts on the topic of active learning, adaptive design, sequential design, and/or Bayesian optimization in materials science. We briefly review these, deferring to the individual publications for details.

One of the earliest works can be traced to Seko et al.^{80,81} who used kriging in conjunction with a simplified version of the probability of improvement to guide DFT calculations towards targeted regions in the design space. They demonstrated this approach to computationally find compounds with high melting points and low lattice thermal conductivity. A similar strategy, using kriging together with simplified probability of improvement, was adopted by Kiyohara et al.⁸² to computationally identify a stable interface structure for a simplified coincidence-site lattice

grain boundary of FCC-Cu. Ueno et al. developed a Bayesian optimization software framework (referred to as COMBO) for black-box optimization, especially for large data problems.³³ The COMBO has several functionalities including Thompson sampling (which is based on probability matching that chooses a candidate point according to the probability of being optimal), one-rank Cholesky update, random feature maps, and automatic hyperparameter tuning. Ling et al.³⁴ emphasize the importance of how estimates for uncertainties are calculated. They bias the usual standard deviation estimate and compare its performance using a RF method on four test data sets including those for magneto-calorics, superconductivity, and thermoelectrics. Comparing strategies in terms of number of evaluations required to find the best property. They find that although random selection was the inferior strategy in all cases, the performance was essentially data dependent. In the limit of adequately large data sets, different uncertainty estimates would be expected to behave similarly, however, an interesting aspect to be explored is the behavior when data sizes are limited, given that we do not know the actual uncertainties associated with the data.⁸³ For Example 1 with the electrostrain data, Xue et al.⁸⁴ have shown that with a data size of 61 the uncertainties calculated by standard deviation and jackknife based on a RF model using 50,000 bootstrap samples, behave very similarly. Yamawaki et al.⁸⁵ use COMBO to perform a structural optimization of graphene nano-ribbons with the objective of identifying a configuration with a high-zT for thermoelectrical properties via separately controlling the electron and phonon transport. Herbol et al.⁸⁶ used Bayesian optimization with expected improvement utility criterion to maximize the intermolecular binding energy for exploring the combinatorial space associated with lead ion solvation in hybrid inorganic–organic perovskites.

The reports discussed above use the formalism of active learning to guide computations or have applied to test cases to demonstrate the potential of these methods in materials science. There have now been a number of studies where iterative machine learning methods have been used to guide experiments. One such work is that of Ren et al.⁸⁷ who used RF to identify glass formers in ternary bulk metallic alloy systems (Co-V-Zr), which in turn were validated by high-throughput combinatorial

experiments. The authors showed the importance of feedback to improve the machine learning model predictions.

Broadening the scope, a number of studies have utilized surrogate based optimization to accelerate materials discovery or codes by minimizing number of experiments or evaluations. Chen et al.⁸⁸ applied the knowledge gradient function with Bayesian inference to study the stability of an emulsion. Its release is triggered by the excitation of gold nanoparticles on the surface of oil droplets and their focus was targeted controlled release through a set of tunable model control parameters. Knowledge gradient was similarly used in maximizing the output current in an optoelectronic device.⁸⁹ Xue et al.²⁹ studied the design of NiTi based alloys with high transformation temperatures. They employed polynomial regression with features for chemical bonding and atomic size coupled with $E[I]$, knowledge gradient and maximum variance to guide the next experiments. Aggarwal et al.²⁵ studied a Cahn-Hilliard model with Bayesian optimization, maximizing the expected utility of the KL divergence to select experiments that provide maximal information about parameters in the model given a certain budget of experiments. Their problem of interest was to find the subsurface properties in a thin film deposited on a substrate. They show that the information gained with two values of the parameters is comparable to that with eight randomly selected values of the same parameters, thus reducing the experimental effort compared to a random strategy by roughly a factor of four. They also studied the problem of metal-metal interfaces and trapping of helium (He) impurities, important in nuclear materials, and also used KL divergence for model selection amongst competing models.

In the field of ferroelectrics, Xue et al. used a Landau model together with Bayesian regression as well as support vector regression combined with $E[I]$ to find Pb-free piezoelectrics in the BaTiO₃ family with morphotropic phase boundaries that show minimum sensitivity to temperature.³⁰ Their training data consisted of 20 well characterized phase diagrams of solid solutions with a search space of 1200 possible phase diagrams. Balachandran et al.³⁶ developed a novel two-step machine learning approach that combines classification learning with regression and adaptive design methods to recommend promising chemical compositions of Bi-containing PbTiO₃ solid solutions that simultaneously satisfy two constraints: (i) They should have a perovskite crystal structure and (ii) They should have a high ferroelectric Curie temperature. From screening a set of 61,000 potential compositions, ten compositions were identified in five iterations and six of them were confirmed to be in the perovskite crystal structure from X-ray diffraction with a composition in the Bi (Co,Fe)O₃-PbTiO₃ chemical space with a highest Curie temperature of 898 K among them.

Optimizing interphase properties in polymer nanocomposites was the problem studied by Wang et al.⁹⁰ Their objective was to minimize the difference between predicted bulk properties of a nanocomposite using simulations for dielectric and viscoelastic properties with those obtained from experiments using a GP model together with $E[I]$.

Genetic algorithms have been extensively used as a means to optimize given target properties iteratively in a variant of the approaches we have considered. Starting with a random initial population of n -block polymers (for any given n), and letting them undergo evolution in terms of constituent blocks and their neighbors based on a genetic algorithm, a set of polymers may be obtained with properties closest to the provided targets.⁹¹ Such methods can be efficient in the search for materials with desired properties compared to random searches, however, they do not use the exploitation—exploration type strategies based on selection and ranking discussed here.

Active learning methods have also been used to parameterize interatomic potentials for molecular dynamics simulations. Podryabinkin and Shapeev⁹² demonstrated the use of active

learning for building machine learning interatomic potentials to resolve the transferability or extrapolation problem. Using a D-optimality criterion, which minimizes the generalized variance of the parameter estimates of a pre-specified model functional form, the authors showed that reliable machine learning interatomic potentials can be built. In a recent paper entitled, “Less is more: sampling chemical space with active learning”. Smith et al.⁹³ also consider the development of accurate and transferrable potentials for predicting molecular energetics. They use a “query by committee” ranking strategy to sample the chemical space where the potential fails to accurately predict the potential energy, thereby improving the overall fitness of the potentials.

Finally, in addition to the use of active learning methods in computational materials and physics which we have discussed in Sec. IVB, there is the whole area of the design of computer experiments,^{94–96} which utilizes adaptive and iterative strategies to address various problems by combining system knowledge with space-filling for sample placement. Many physical processes are difficult to simulate and a simulation code often serves as a proxy for the process with various inputs and outputs. Our Example 2 was a case in point and, with easier access to more powerful computer resources, such studies are becoming important as surrogates to complement physical experiments. Computer experiments allow many of the criteria and designs, such as two-level factorials, Latin hypercubes, U- designs, lattice designs to be compared. In addition, in purely deterministic computer experiments, the observations from the code with the same input variables will be the same. Thus, in modeling data from such an experiment, bias due to model inadequacy becomes important rather than the need to reduce variance.

CHALLENGES AND FUTURE DEVELOPMENT

The nature of experimental data in materials science is such that we do not have the luxury to randomly select different training sets from a population and evaluate the ground truth, because each experiment is expensive. Typically we have a training data set of samples from an unknown population and we need to employ statistical methods, such as non-parametric bootstrap sampling with replacement (and other variants of this methodology) which we have discussed, to generate many training data sets to form a prescribed distribution to train surrogate models.

There are different ways to evaluate the efficacy of AL methods. One of the approaches is to use simulated datasets, where the ground truth labels are known for all data points in the population. This approach has advantages because one can randomly select several training sets from a population, explore AL methods independently on each of the training set, and provide statistically meaningful comparisons. In fact, such studies on different AL strategies on simulated datasets have been performed in the literature.^{97,98} For example, Theiler and Zimmer⁹⁷ compare Bayesian global optimization methods using simulated datasets in four distinct regimes, corresponding to high and low levels of measurement noise and to high and low levels of “quenched noise.” They investigate the choice of AL methods when the surrogate model is well matched to the data. Their numerical experiments provide a direct way to compare Bayesian global optimization algorithms, and in particular, in a way that controls for the choice of surrogate model and avoids the issue of model mismatch. They isolate the role of the surrogate model from the AL. So, rather than choosing anecdotal functions to optimize, they draw their functions (thousands of them) from a GP prior, and then use the same prior to initialize the GP regression. By comparing these algorithms over different regimes of quenched and measurement noise, they compare relative performance of the different AL strategies. Their work on gaussian generated data categorically demonstrates that random sampling is the worst performer compared to other strategies. Similarly,

Balachandran et al.⁹⁹ used data from DFT calculations and evaluated various AL methods with the objective of accelerated materials design, albeit *post factum*. Here, the authors considered a population of 223 compounds for which the ground truth data was available. They were able to randomly sample a number of training sets from the population and estimate the statistics. Even in this case, AL consistently out-performed random sampling.

One of the unresolved questions in adaptive design is the importance of model quality in the accelerated search. We have empirically observed that it is not necessary to have a highly accurate machine learned model to accelerate the search when active learning is also involved in recommending the next experiment or computation. In Fig. 12a, we show a particular example of a machine learned model that was trained using the DFT data. The term “G” refers to shear modulus of a particular compound⁹⁹ obtained from DFT calculations. It can be seen that this is particularly not an accurate model. However, when this model is coupled to active learning algorithms (E[I], KG etc) in search of a new compound with the largest G, then it performs better than the random method (see Fig. 12b). This result indicates that the active learning is forgiving of the poor model quality. Our preliminary studies have shown that the choice of features may also have a key role in finding the optimal material when dealing with poor model quality. However, the underlying reasons behind these observations are not clear and worthy of further investigation.

Similarly, there is a need for results providing stricter guidance on stopping criteria to prevent excessive exploration of the feature space. In the examples we have considered, we have stopped when either our budget has run out or when we are satisfied with the outcome.

The methods we have outlined have been generalized to deal with multi-objectives so that given two properties, one needs to make the best selection for the next materials in parts of the two dimensional Pareto plot likely to improve on the existing Pareto front.^{100–103} The Pareto plot has the axes as the properties so that we can define a characteristic boundary on which lie materials where none of the objectives can be improved in value without degrading the other objective value. Such boundary points, the non-dominated data-points, define a Pareto front (PF) that represents the best trade-off between the objectives. Although the method has been exercised on materials databases,¹⁰³ such as for shape memory alloys and piezoelectrics, it yet has to be demonstrated experimentally to find new materials with enhanced properties.

The extensions of the basic ideas here to multi-fidelity optimization, in which larger amounts of relatively cheap data are used in conjunction with smaller amounts of more expensive data to make credible predictions, has also been well studied. These methods employ GP and extend well known kriging regression to heterogeneous variable-fidelity information,^{104,105} and this framework has been successfully demonstrated within an optimization scheme. Its use in materials science has recently been demonstrated by combining two levels of DFT exchange correlation functionals (PBE and Hybrid functionals) to make predictions of band gaps for double perovskites.¹⁰⁶ However, no design was performed and the method still needs to be demonstrated with experiments. Nevertheless, there are a number of materials science problems where this method will become important.

Recently, there has been interest in the use of text mining methods and natural language processing techniques to build datasets from the digital literature for guiding materials synthesis and design.^{107,108} Such advances can have key implications with the adaptive learning approach, especially in rapidly constructing the training datasets for building surrogate models. Currently the training sets are constructed manually from surveying the published literature, which can be a time-consuming process.

We envision that combining the text mining approach for training set construction with adaptive learning can have an impact in the accelerated search for new materials.

SUMMARY

Our objective is to bring global optimization based tools, developed in other areas, to the attention of the materials community by showing via three diverse examples that involve experiments or computations how the number of experiments or iterations can be minimized. Although a large number of studies exist that build and make predictions using machine learning models, the importance of experimental design and active learning—the selection of the next material to test—is only now being studied in materials science. In particular, we have highlighted the importance of how different design strategies, in the form of surrogate models and utility functions, can minimize the number of iterations to find materials with desired properties. These tools have certainly led to new materials in their class and other aspects as illustrated by our three examples, and the variety of other examples we have briefly discussed in Sec. IVC. Other applications relate to processing,³⁹ as well as the use of computational codes for design in the materials and mechanical engineering areas^{14,37} and functional polymers.⁹¹ Many of these studies (experimental and computational) have shown that the design strategies are superior to random selection in minimizing the number of experiments. We also note as demonstrated by Example 3 in the context of MAX and apatite compounds, that the surrogate model does not necessarily improve over the course of the iterations, even though it leads to fewer measurements or better properties in the course of the iterations (Fig. 12a, b). This aspect has been observed also in other studies.^{27,30} There is thus a need for statistical studies of the interplay between the quality of the model and the selections recommended by the design strategies. However, we need to be cognizant of the “no free lunch theorem,”¹⁰⁹ which essentially states that there is no universal optimizer for all problems! That is, there is no guarantee that a model trained on one data set will work on a different data set.

We are still at the stage where these methods have been barely exercised experimentally. The examples we have given have involved bulk synthesis with relatively slow turn around time and characterization. High throughput experiments performing synthesis by sputtering and X-ray have had success with thin film geometries¹¹⁰ as screening strategies for shape memory alloy design in bulk. For bulk synthesis, self-propagation high-temperature synthesis (SHS),¹¹¹ methods based on laser additive manufacturing,¹¹² and suspended droplet alloy processing show significant promise. The holy grail of automated on the fly synthesis and characterization will require the kinds of design tools we have discussed for next material selection. Recent steps in this direction include the use of unsupervised learning, such as cluster analysis, to convert composition and structural data from combinatorial libraries and material structure databases into potential composition phase diagrams and potential constituent phases.^{113,114} The convergence of the information sciences with materials science augurs considerable promise in changing the way new materials with targeted properties will be discovered, and how materials science will be studied in the future.

DATA AVAILABILITY

The dataset for Example 1, the electrostrain problem, is included in the Supplemental file.

ACKNOWLEDGEMENTS

We acknowledge many useful discussions with our colleagues Jim Gubernatis, John Hogden, James Theiler, Bertrand Leduc on aspects of the work reviewed here. We are

grateful to the Laboratory Directed Research and Development (LDRD) program (project #20180660ER) and the Center for Nonlinear Studies at Los Alamos National Laboratory for support.

AUTHOR CONTRIBUTIONS

All authors have contributed to the writing of this manuscript.

ADDITIONAL INFORMATION

Supplementary information accompanies the paper on the *npj Computational Materials* website (<https://doi.org/10.1038/s41524-019-0153-8>).

Competing interests: The authors declare no competing interests.

Publisher's note: Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

REFERENCES

- OSTP. *Materials Genome Initiative for Global Competitiveness* (https://www.whitehouse.gov/sites/default/files/microsites/ostp/materials_genome_initiative_final.pdf) (2011).
- Curtarolo, S. et al. AFLOWLIB.ORG: a distributed materials property repository from high-throughput *ab initio* calculations. *Comput. Mater. Sci.* **58**, 227–235 (2012).
- Jain, A. et al. Commentary: The Materials Project: a materials genome approach to accelerating materials innovation. *APL Mater.* **1** (2013).
- Saal, J. E., Kirklin, S., Aykol, M., Meredig, B. & Wolverton, C. Materials design and discovery with high-throughput density functional theory: The Open Quantum Materials Database (OQMD). *JOM* **65**, 1501–1509 (2013).
- Green, M. L. et al. Fulfilling the promise of the materials genome initiative with high-throughput experimental methodologies. *Appl. Phys. Rev.* **4**, 011105 (2017).
- Atkinson, A. & Donev, A. *Optimum Experimental Designs*. (Oxford University Press, Oxford, 1992).
- Park, J.-M. & Hu, Y. H. On-line learning for active pattern recognition. *IEEE Signal Process. Lett.* **3**, 301–303 (1996).
- Menache, I., Mannor, S. & Shimkin, N. Basis function adaptation in temporal difference reinforcement learning. *Ann. Oper. Res.* **134**, 215–238 (2005).
- Kaelbling, L. *Learning in Embedded Systems*. (MIT Press, Cambridge, MA, 1993).
- Murphy, R. F. An active role for machine learning in drug development. *Nat. Chem. Biol.* **7**, 327 (2011).
- Mohamed, T. P., Carbonell, J. G. & Ganapathiraju, M. K. Active learning for human protein–protein interaction prediction. *BMC Bioinform.* **11**, S57 (2010).
- Stegle, O., Payet, L., Mergny, J.-L., MacKay, D. J. C. & Huppert, J. L. Predicting and understanding the stability of G-quadruplexes. *Bioinformatics* **25**, i374–i1382 (2009).
- Liu, Y. Active learning with support vector machine applied to gene expression data for cancer classification. *J. Chem. Inf. Comput. Sci.* **44**, 1936–1941 (2004).
- Forrester, A. I. J., Sobester, A. & Keane, A. J. *Engineering Design via Surrogate Modelling: A Practical Guide* (Wiley, Hoboken, 2008).
- Nishijima, M. et al. Accelerated discovery of cathode materials with prolonged cycle life for lithium-ion battery. *Nat. Commun.* **5**, 4553 (2014).
- Sanvito, S. et al. Accelerated discovery of new magnets in the Heusler alloy family. *Sci. Adv.* **3** (2017).
- Sharma, V. et al. Rational design of all organic polymer dielectrics. *Nat. Commun.* **5**, 4845 (2014).
- Sparks, T. D., Gaultois, M. W., Oliynyk, A., Brgoch, J. & Meredig, B. Data mining our way to the next generation of thermoelectrics. *Scr. Mater.* **111**, 10–15 (2016).
- Rajan, K. Materials informatics: the materials “gene” and big data. *Annu. Rev. Mater. Res.* **45**, 153–169 (2015).
- Kalidindi, S. R. & Graef, M. D. Materials data science: current status and future outlook. *Annu. Rev. Mater. Res.* **45**, 171–193 (2015).
- Ramprasad, R., Batra, R., Piliand, G., Mannodi-Kanakithodi, A. & Kim, C. Machine learning in materials informatics: recent applications and prospects. *npj Comput. Mater.* **3**, 54 (2017).
- Hill, J. et al. Materials science with large-scale data and informatics: unlocking new opportunities. *MRS Bull.* **41**, 399–409 (2016).
- Ward, L. & Wolverton, C. Atomistic calculations and materials informatics: a review. *Curr. Opin. Solid State Mater. Sci.* **21**, 167–176 (2017).
- Lookman, T., Balachandran, P. V., Xue, D., Hogden, J. & Theiler, J. Statistical inference and adaptive design for materials discovery. *Curr. Opin. Solid State Mater. Sci.* **21**, 121–128 (2017).
- Aggarwal, R., Demkowicz, M. J. & Marzouk, Y. M. *Information-Driven Experimental Design in Materials Science*, 13–44 (Springer, New York, 2016).
- Balachandran, P. V., Young, J., Lookman, T. & Rondinelli, J. M. Learning from data to design functional materials without inversion symmetry. *Nat. Commun.* **8**, 14282 (2017).
- Xue, D. et al. Accelerated search for materials with targeted properties by adaptive design. *Nat. Commun.* **7**, 11241 (2016).
- Dehghannasiri, R. et al. Optimal experimental design for materials discovery. *Comput. Mater. Sci.* **129**, 311–322 (2017).
- Xue, D. et al. An informatics approach to transformation temperatures of NiTi-based shape memory alloys. *Acta Mater.* **125**, 532–541 (2017).
- Xue, D. et al. Accelerated search for BaTiO₃-based piezoelectrics with vertical morphotropic phase boundary using Bayesian learning. *Proc. Natl Acad. Sci.* **113**, 13301–13306 (2016).
- Rouet-Leduc, B., Barros, K., Lookman, T. & Humphreys, C. J. Optimisation of gan leds and the reduction of efficiency droop using active machine learning. *Sci. Rep.* **6**, 24862 (2016).
- Rouet-Leduc, B., Hulbert, C., Barros, K., Lookman, T. & Humphreys, C. J. Automated convergence of optoelectronic simulations using active machine learning. *Appl. Phys. Lett.* **111**, 043506 (2017).
- Ueno, T., Rhone, T. D., Hou, Z., Mizoguchi, T. & Tsuda, K. COMBO: an efficient Bayesian optimization library for materials science. *Mater. Discov.* **4**, 18–21 (2016).
- Ling, J., Hutchinson, M., Antono, E., Paradiso, S. & Meredig, B. High-dimensional materials and process optimization using data-driven experimental design with well-calibrated uncertainty estimates. *Integr. Mater. Manuf. Innov.* **6**, 207–217 (2017).
- Yuan, R. et al. Accelerated discovery of large electrostrains in BaTiO₃-based piezoelectrics using active learning. *Adv. Mater.* 1702884 (2018).
- Balachandran, P. V., Kowalski, B., Sehirlioglu, A. & Lookman, T. Experimental search for high-temperature ferroelectric perovskites guided by two-step machine learning. *Nat. Commun.* **9**, 1668 (2018).
- McDowell, D. L. et al. *Integrated Design of Multiscale, Multifunctional Materials and Products* (Butterworth-Heinemann, 2009).
- Jones, D. R., Schonlau, M. & Welch, W. J. Efficient global optimization of expensive black-box functions. *J. Glob. Optim.* **13**, 455–492 (1998).
- Casciato, M. J., Kim, S., Lu, J. C., Hess, D. W. & Grover, M. A. Optimization of a carbon dioxide-assisted nanoparticle deposition process using sequential experimental design with adaptive design space. *Ind. Eng. Chem. Res.* **51**, 4363–4370 (2012).
- Brochu, E., Cora, V. M. & de Freitas, N. A tutorial on Bayesian optimization of expensive cost functions, with application to active user modeling and hierarchical reinforcement learning. *ArXiv e-prints* (2010).
- Chaloner, K. & Verdine, I. Bayesian experimental design: a review. *Stat. Sci.* **10**, 273–304 (1995).
- Lindley, D. V. On a measure of the information provided by an experiment. *Ann. Math. Stat.* **27**, 986–1005 (1956).
- Frazier, P., Powell, W. & Dayanik, S. The knowledge gradient policy for correlated normal beliefs. *INFORMS J. Comput.* **21**, 599–613 (2009).
- Huan, X. & Marzouk, Y. M. Simulation-based optimal Bayesian experimental design for nonlinear systems. *J. Comput. Phys.* **232**, 288–317 (2013).
- Cavazzuti, M. *Optimization Methods*. (Springer, Berlin, Heidelberg, 2013).
- Fisher, R. *The Design of Experiments*, 5th ed. (Oliver & Boyd, Oxford, 1949).
- Box, G. E. P. & Wilson, K. B. On the experimental attainment of optimum conditions. *J. R. Stat. Soc. Ser. B (Methodol.)* **13**, 1–45 (1951).
- Kackar, R. N. Off-line quality control, parameter design, and the Taguchi method. *J. Qual. Technol.* **17**, 176–188 (1985).
- Kushner, H. A new method of locating the maximum of an arbitrary multi-peak curve in the presence of noise. *J. Basic Eng.* **86**, 97–106 (1964).
- Mockus, J., Tiesis, V. & Zilinskas, A. in *Towards Global Optimisation*, Vol. 2 (eds Dixon, L. & Szego, G.) 117–129 (Elsevier Science Ltd., North Holland, Amsterdam, 1978).
- Snoek, J., Larochelle, H. & Adams, R. Practical bayesian optimization of machine learning algorithms. *Adv. Neural Inf. Process. Syst.* 2951–2959 (2012).
- Howard, R. Information value theory. *IEEE Trans. Syst. Sci. Cybern.* **2**, 22–26 (1966).
- Lindley, D. V. & Smith, A. F. M. Bayes estimates for the linear model. *J. R. Stat. Soc. Ser. B* **34**, 1–41 (1972).
- Clyde, M. A. Experimental design: a Bayesian perspective (2007).
- O'Hagan, A. & Kingman, J. F. C. Curve fitting and optimal design for prediction. *J. R. Stat. Soc. Ser. B* **40**, 1–42 (1978).
- Mockus, J. Application of Bayesian approach to numerical methods of global and stochastic optimization. *J. Glob. Optim.* **4**, 347–365 (1994).
- Kaelbling, L. P., Littman, M. L. & Moore, A. W. Reinforcement learning: a survey. *J. Artif. Intell. Res.* **4**, 237–285 (1996).

58. Sutton, R. S. Learning to predict by the methods of temporal differences. *Mach. Learn.* **3**, 9–44 (1988).
59. Powell, W. B. & Ryzhov, I. O. *Optimal Learning* (Wiley, Hoboken, 2012).
60. Scott, W., Frazier, P. I. & Powell, W. The correlated knowledge gradient for simulation optimization of continuous parameters using gaussian process regression. *SIAM J. Optim.* **21**, 996–1026 (2011).
61. Letham, B., Karrery, B., Ottoniz, G. & Bakshy, E. Constrained Bayesian optimization with noisy experiments [abs/1706.07094v1](https://arxiv.org/abs/1706.07094v1) (2017).
62. Yoon, B. J., Qian, X. & Dougherty, E. R. Quantifying the objective cost of uncertainty in complex dynamical systems. *IEEE Trans. Signal Process.* **61**, 2256–2266 (2013).
63. Dehghannasiri, R., Yoon, B.-J. & Dougherty, E. R. Optimal experimental design for gene regulatory networks in the presence of uncertainty. *IEEE/ACM Trans. Comput. Biol. Bioinform.* **12**, 938–950 (2015).
64. Dehghannasiri, R., Yoon, B.-J. & Dougherty, E. R. Efficient experimental design for uncertainty reduction in gene regulatory networks. *BMC Bioinform.* **16**, S2 (2015).
65. Efron, B. & Gong, G. A leisurely look at the bootstrap, the jackknife, and cross-validation. *Am. Stat.* **37**, 36–48 (1983).
66. Piprek, J. *Nitride Semiconductor Devices: Principles and Simulation* (Wiley, Hoboken, 2007).
67. Wild, S. M., Sarich, J. & Schunck, N. Derivative-free optimization for parameter estimation in computational nuclear physics. *J. Phys. G* **42**, 034031 (2015).
68. Pang, G., Perdikaris, P., Cai, W. & Karniadakis, G. E. Discovering variable fractional orders of advection–dispersion equations from field data using multi-fidelity bayesian optimization. *J. Comput. Phys.* **348**, 694–714 (2017).
69. Gosavi, A. *Simulation-based Optimization: Parametric Optimization Techniques and Reinforcement Learning*. (Operations Research/Computer Science Interfaces Series, ISBN-13: 978-1489974907, Second Edition 2015).
70. Sacks, J., Welch, W. J., Mitchell, T. J. & Wynn, H. P. Design and analysis of computer experiments. *Stat. Sci.* **4**, 1–11 (2018).
71. Hutter, F., Hoos, H. H. & Leyton-Brown, K. Sequential model-based optimization for general algorithm configuration. in *International Conference on Learning and Intelligent Optimization*, 507–523 (Springer, 2011).
72. Bergstra, J. S., Bardenet, R., Bengio, Y. & Kégl, B. Algorithms for hyper-parameter optimization. in *Advances in Neural Information Processing Systems*, 2546–2554 (2011).
73. de Jong, M. et al. A statistical learning framework for materials science: application to elastic moduli of k-nary inorganic polycrystalline compounds. *Sci. Rep.* **6**, 34256 (2016).
74. Ward, L., Agrawal, A., Choudhary, A. & Wolverton, C. A general-purpose machine learning framework for predicting properties of inorganic materials. *npj Comput. Mater.* **2**, 16028 (2016).
75. Seko, A., Hayashi, H., Nakayama, K., Takahashi, A. & Tanaka, I. Representation of compounds for machine-learning prediction of physical properties. *Phys. Rev. B* **95**, 144110 (2017).
76. Balachandran, P. V., Shearman, T., Theiler, J. & Lookman, T. Predicting displacements of octahedral cations in ferroelectric perovskites using machine learning. *Acta Crystallogr. Sect. B* **73**, 962–967 (2017).
77. Balachandran, P. V. et al. in *Materials Discovery and Design*, 59–79 (Springer, New York, 2018).
78. White, T., Ferraris, C., Kim, J. & Madhavi, S. Apatite—an adaptive framework structure. *Rev. Mineral. Geochem.* **57**, 307–401 (2005).
79. Shannon, R. D. Revised effective ionic radii and systematic studies of interatomic distances in halides and chalcogenides. *Acta Cryst. A* **32**, 751–767 (1976).
80. Seko, A., Maekawa, T., Tsuda, K. & Tanaka, I. Machine learning with systematic density-functional theory calculations: Application to melting temperatures of single- and binary-component solids. *Phys. Rev. B* **89**, 054303 (2014).
81. Seko, A. et al. Prediction of low-thermal-conductivity compounds with first-principles anharmonic lattice-dynamics calculations and bayesian optimization. *Phys. Rev. Lett.* **115**, 205901 (2015).
82. Kiyohara, S., Oda, H., Tsuda, K. & Mizoguchi, T. Acceleration of stable interface structure searching using a kriging approach. *Jpn. J. Appl. Phys.* **55**, 045502 (2016).
83. Wager, S., Hastie, T. & Efron, B. Confidence intervals for random forests: the jackknife and the infinitesimal jackknife. *J. Mach. Learn. Res.* **15**, 1625–1651 (2014).
84. Xue, D., Tian, Y., Yuan, R. & Lookman, T. Bayesian optimization in designing shape memory alloy. (eds. Foiles, S., McDowell, D., Strachan, A.). *Model. Simul. Mater. Sci. Eng.* (2019).
85. Yamawaki, M., Ohnishi, M., Ju, S. & Shiomi, J. Multifunctional structural design of graphene thermoelectrics by Bayesian optimization. *Sci. Adv.* **4**, 1–8 (2018).
86. Herbol, H. C., Hu, W., Frazier, P., Clancy, P. & Poloczek, M. Efficient search of compositional space for hybrid organic–inorganic perovskites via Bayesian optimization. *npj Comput. Mater.* **4**, 51 (2018).
87. Ren, F. et al. Accelerated discovery of metallic glasses through iteration of machine learning and high-throughput experiments. *Sci. Adv.* **4**, 1–11 (2018).
88. Chen, S., Reyes, K., Gupta, M., McAlpine, M. & Powell, W. Optimal learning in experimental design using the knowledge gradient policy with application to characterizing nanoemulsion stability. *SIAM/ASA J. Uncertainty Quant.* **3**, 320–345 (2015).
89. Wang, Y., Reyes, K., Brown, K., Mirkin, C. & Powell, W. Nested-batch-mode learning and stochastic optimization with an application to sequential multi-stage testing in materials science. *SIAM J. Sci. Comput.* **37**, B361–B381 (2015).
90. Wang, Y. et al. Identifying interphase properties in polymer nanocomposites using adaptive optimization. *Compos. Sci. Technol.* **162**, 146–155 (2018).
91. Mannodi-Kanakithodi, A., Piliñia, G., Huan, T. D., Lookman, T. & Ramprasad, R. Machine learning strategy for accelerated design of polymer dielectrics. *Sci. Rep.* **6**, 20952 EP (2016).
92. Podryabinkin, E. V. & Shapeev, A. V. Active learning of linearly parametrized interatomic potentials. *Comput. Mater. Sci.* **140**, 171–180 (2017).
93. Smith, J. S., Nebgen, B., Lubbers, N., Isayev, O. & Roitberg, A. E. Less is more: sampling chemical space with active learning. *J. Chem. Phys.* **148**, 241733 (2018).
94. Garud, S. S., Karimi, I. A. & Kraft, M. Design of computer experiments: a review. *Comput. Chem. Eng.* **106**, 71–95 (2017).
95. Santner, T. J., Williams, B. J. & Notz, W. I. *The Design and Analysis of Computer Experiments* (Springer, New York, 2018, ISBN: 978-1-4939-8845-7, first edition 2003).
96. Bursztyn, D. & Steinberg, D. M. Comparison of designs for computer experiments. *J. Stat. Plan. Inference* **136**, 1103–1119 (2006).
97. Theiler, J. & Zimmer, B. G. Selecting the selector: comparison of update rules for discrete global optimization. *Stat. Anal. Data Min.* **10**, 211–229 (2017).
98. Picheny, V., Ginsbourger, D., Richet, Y. & Caplin, G. Quantile-based optimization of noisy computer experiments with tunable precision. *Technometrics* **55**, 2–13 (2013).
99. Balachandran, P. V., Xue, D., Theiler, J., Hogden, J. & Lookman, T. Adaptive strategies for materials design using uncertainties. *Sci. Rep.* **6**, 19660 (2016).
100. Keane, A. J. Statistical improvement criteria for use in multiobjective design optimization. *AIAA J.* **44**, 879–891 (2006).
101. Svenson, J. & Santner, T. Multiobjective optimization of expensive-to-evaluate deterministic computer simulator models. *Comput. Stat. Data Anal.* **94**, 250–264 (2016).
102. Wagner, T., Emmerich, M., Deutz, A. & Ponweiser, W. in *Parallel Problem Solving from Nature, PPSN XI*, (eds. Schaefer, R., Cotta, C., Kołodziej, J. & Rudolph, G.) 718–727 (Springer, Berlin Heidelberg, 2010).
103. Gopakumar, A. M., Balachandran, P. V., Xue, D., Gubernatis, J. E. & Lookman, T. Multi-objective optimization for materials discovery via adaptive design. *Sci. Rep.* **8**, 3738 (2018).
104. Kennedy, M. C. & O’Hagan, A. Predicting the output from a complex computer code when fast approximations are available. *Biometrika* **87**, 1–13 (2000).
105. Forrester, A. I., Sobester, A. & Keane, A. J. Multi-fidelity optimization via surrogate modelling. *Proc. R. Soc. Lond. A* **463**, 3251–3269 (2007).
106. Piliñia, G., Gubernatis, J. & Lookman, T. Multi-fidelity machine learning models for accurate bandgap predictions of solids. *Comput. Mater. Sci.* **129**, 156–163 (2017).
107. Kim, E. et al. Materials synthesis insights from scientific literature via text extraction and machine learning. *Chem. Mater.* **29**, 9436–9444 (2017).
108. Kim, E. et al. Machine-learned and codified synthesis parameters of oxide materials. *Sci. Data* **4**, 170127 (2017).
109. Wolpert, D. H. & Macready, W. G. No free lunch theorems for optimization. *IEEE Trans. Evol. Comput.* **1**, 67–82 (1997).
110. Koinuma, H. & Takeuchi, I. Combinatorial solid-state chemistry of inorganic materials. *Nat. Mater.* **3**, 429 (2004).
111. Resnina, N. et al. Martensitic transformation and mechanical behavior of porous Ti-50.0 at% Ni alloy, fabricated by self-propagating high temperature synthesis at different temperature. *Phys. Procedia* **10**, 11–16 (2010).
112. Polanski, M., Kwiatkowska, M., Kunce, I. & Bystrzycki, J. Combinatorial synthesis of alloy libraries with a progressive composition gradient using laser engineered net shaping (LENS): hydrogen storage alloys. *Int. J. Hydrog. Energy* **38**, 12159–12171 (2013).
113. Kusne, A. G., Keller, D., Anderson, A., Zaban, A. & Takeuchi, I. High-throughput determination of structural phase diagram and constituent phases using GRENDEL. *Nanotechnology* **26**, 444002 (2015).
114. Kusne, A. G. et al. On-the-fly machine-learning for high-throughput experiments: search for rare-earth-free permanent magnets. *Sci. Rep.* **4**, 6367 (2014).



Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons license, unless indicated otherwise in a credit line to the material. If material is not included in the

article's Creative Commons license and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this license, visit <http://creativecommons.org/licenses/by/4.0/>.

© The Author(s) 2019