# **Computational catalyst discovery: Active** classification through myopic multiscale sampling

Cite as: J. Chem. Phys. 154, 124118 (2021); https://doi.org/10.1063/5.0044989 Submitted: 21 January 2021 . Accepted: 09 March 2021 . Published Online: 25 March 2021

២ Kevin Tran, ២ Willie Neiswanger, Kirby Broderick, Eric Xing, Jeff Schneider, and ២ Zachary W. Ulissi

# COLLECTIONS

Paper published as part of the special topic on Computational Materials Discovery

# **ARTICLES YOU MAY BE INTERESTED IN**

Uncertainty estimation for molecular dynamics and sampling The Journal of Chemical Physics 154, 074102 (2021); https://doi.org/10.1063/5.0036522

Analytical gradients for molecular-orbital-based machine learning The Journal of Chemical Physics 154, 124120 (2021); https://doi.org/10.1063/5.0040782

Improving molecular force fields across configurational space by combining supervised and unsupervised machine learning

The Journal of Chemical Physics 154, 124102 (2021); https://doi.org/10.1063/5.0035530





J. Chem. Phys. 154, 124118 (2021); https://doi.org/10.1063/5.0044989 © 2021 Author(s).

Export Citation

# Computational catalyst discovery: Active classification through myopic multiscale sampling

Cite as: J. Chem. Phys. 154, 124118 (2021); doi: 10.1063/5.0044989 Submitted: 21 January 2021 • Accepted: 9 March 2021 • Published Online: 25 March 2021

Kevin Tran,<sup>1</sup> D Willie Neiswanger,<sup>2</sup> Kirby Broderick,<sup>1</sup> Eric Xing,<sup>23,4</sup> Jeff Schneider,<sup>2</sup> and Zachary W. Ulissi<sup>1,a)</sup>

## AFFILIATIONS

<sup>1</sup>Chemical Engineering Department, Carnegie Mellon University, Pittsburgh, Pennsylvania 15217, USA
 <sup>2</sup>Machine Learning Department, Carnegie Mellon University, Pittsburgh, Pennsylvania 15217, USA
 <sup>3</sup>Mohamed bin Zayed University of Artificial Intelligence, Masdar City, Abu Dhabi, United Arab Emirates
 <sup>4</sup>Petuum Inc., Pittsburgh, Pennsylvania 15213, USA

**Note:** This paper is part of the JCP Special Topic on Computational Materials Discovery. <sup>a)</sup>**Author to whom correspondence should be addressed:** zulissi@andrew.cmu.edu

## ABSTRACT

The recent boom in computational chemistry has enabled several projects aimed at discovering useful materials or catalysts. We acknowledge and address two recurring issues in the field of computational catalyst discovery. First, calculating macro-scale catalyst properties is not straightforward when using ensembles of atomic-scale calculations [e.g., density functional theory (DFT)]. We attempt to address this issue by creating a multi-scale model that estimates bulk catalyst activity using adsorption energy predictions from both DFT and machine learning models. The second issue is that many catalyst discovery efforts seek to optimize catalyst properties, but optimization is an inherently exploitative objective that is in tension with the explorative nature of early-stage discovery projects. In other words, why invest so much time finding a "best" catalyst when it is likely to fail for some other, unforeseen problem? We address this issue by relaxing the catalyst discovery goal into a classification problem: "What is the set of catalysts that is worth testing experimentally?" Here, we present a catalyst discovery method called myopic multiscale sampling, which combines multiscale modeling with automated selection of DFT calculations. It is an active classification strategy that seeks to classify catalysts as "worth investigating" or "not worth investigating" experimentally. Our results show an ~7–16 times speedup in catalyst classification relative to random sampling. These results were based on offline simulations of our algorithm on two different datasets: a larger, synthesized dataset and a smaller, real dataset.

Published under license by AIP Publishing. https://doi.org/10.1063/5.0044989

# I. INTRODUCTION

Recent advances in computing hardware and software have led to substantial growth in the field of computational materials science. In particular, databases of high-throughput calculations<sup>1–6</sup> have increased the amount of information available to researchers. These databases facilitate the development of models that supplement human understanding of physical trends in materials.<sup>7–9</sup> These models can then be used in experimental discovery efforts by identifying promising subsets of the search space, resulting in increased experimental efficiency.<sup>10–15</sup>

However, many materials design efforts use material properties and calculation archetypes that are too problem-specific to be tabulated in generalized databases. When such efforts coincide with design spaces too large to search in a feasible amount of time, we need a way to search through the design space efficiently. Sequential learning, sometimes referred to as optimal design of experiments or active learning, can fill this role. Sequential learning is the process of using the currently available data to decide which new data would be most valuable for achieving a particular goal.<sup>16–18</sup> In practice, this usually involves fitting a surrogate model to the available data and then pairing the model with an *acquisition function* that calculates the values of a new, potential data point. Then, we *query* the most valuable data points, add them to the dataset, and repeat this process. These sequential learning methods have been estimated to accelerate materials discovery efforts by up to a factor of 20.<sup>19</sup>

Sequential learning has numerous sub-types of methods that can and have been used for different goals. One such sub-type is

active learning. With many active learning algorithms, the goal is to replace a relatively slow data-querying process with a faster-running surrogate model.<sup>20</sup> Since the surrogate model may be used to query any point, the acquisition functions focus on ensuring that the entire search space is explored. Another sub-type of sequential learning is active optimization.<sup>21</sup> With this sub-type, the goal is to maximize or minimize some objective function. Thus, the acquisition functions generally focus on parts of the search space where maxima or minima are more likely to occur. One of the most common types of active optimization is Bayesian optimization.<sup>21</sup> Yet another sub-type of sequential learning is online or on-the-fly learning.<sup>22</sup> The goal of these methods is to accelerate the predictions of streams of data. In the field of computational materials science, this is often applied to predicting trajectories for density functional theory (DFT) or molecular dynamics calculations.<sup>23,24</sup>

In computational materials discovery, we often have the following task: we have a set of available materials  $\mathcal{X} = \{x_i\}_{i=1}^n$ , where each material  $x_i$  has an associated quantity  $y_i$ , denoting its value for some application. Examples of common properties for  $y_i$  include—but are not limited to—formation energies of materials, catalyst activity, tensile strength, or conductivity. The value  $y_i$  is unknown and must be calculated, which can be costly in time, money, or other resources. Furthermore, theoretical calculations of material properties may be inconsistent with experimental results. As per a common aphorism among statisticians, "All models are wrong, but some are useful."

Due to these potential model errors and due to the exploratory nature of materials discovery, we propose reframing the materials discovery question. Instead of trying to discover materials with optimal  $y_i$  values, what if we instead classify materials as having promising or unpromising  $y_i$  values? In other words, what if we frame materials discovery efforts as classification problems rather than optimization problems? The estimated classes could then be used to design physical experiments. Mathematically, this is akin to assuming that the material *i* has a binary value  $y_i \in \{0, 1\}$ , where 0 denotes "*not of interest*" and 1 denotes "*of interest*."

The goal is then to determine the values  $y_i$  for each  $x_i \in \mathcal{X}$  as cheaply as possible. One can view this as the task of most-efficiently learning a classifier that, for each  $x_i$ , correctly predicts its value  $y_i$ . In this way, materials discovery problems can be framed as problems of *active classification*. Active classification is the task of choosing an ordering of  $x_i \in \mathcal{X}$ , over which we will iterate and sequentially measure their values  $y_i$ , in order to most efficiently (using the fewest measurements) learn a classifier that predicts the correct label for all materials  $x_i \in \mathcal{X}$ .<sup>25,26</sup>

Another aspect of computational materials discovery is the ability to turn calculations into recommendations—e.g., how can we convert DFT results into actionable experiments? This conversion is relatively straight-forward when properties are directly calculable, which is the case for properties such as the enthalpy of formation.<sup>27</sup> If we perform a single DFT calculation that suggests a single material may be stable, then we can suggest that single material for experimentation. However, for many applications, the properties of interest may not be calculable directly. For example, let us say we are interested in finding active catalysts. One way to do this is to use DFT to calculate the adsorption energy between the catalyst and particular reaction intermediates and then couple the resulting adsorption energy with a Sabatier relationship.<sup>28</sup> However, *in situ*, a catalyst comprises numerous adsorption sites and surfaces. Thus, the true activity of a catalyst may be governed by an ensemble of adsorption energies and, therefore, may need multiple DFT calculations. How do we address the fact that we need multiple DFT queries to resolve the properties of a single material?

Here, we attempt to address both outlined issues: (1) we need an ensemble of DFT queries to calculate a single experimental property of interest and (2) we need a sequential learning method designed for high-throughput discovery/classification. We overcome both issues by creating the Myopic Multiscale Sampling (MMS) method (Fig. 1). MMS addresses the first aforementioned issue by using a multiscale modeling framework for estimating the activity of a catalyst using an ensemble of both DFT and Machine Learning (ML) predicted adsorption energies. MMS then addresses the second issue by combining this multiscale modeling framework with a number of sequential learning methods, including active classification. Note that MMS, as we describe it in this paper, is tailored to discovering active catalysts. Although this method may not be directly transferable to other applications, we hope that others may be able to adapt the principles of the method to their own applications.

#### **II. METHODS**

#### A. Multiscale modeling

In this paper, we use the discovery of active catalysts as a case study. Catalyst activity is often correlated with the adsorption energy of particular reaction intermediates, as per the volcano relationships stemming from the Sabatier principle.<sup>28,29</sup> These adsorption energies can be calculated using DFT. Each DFT-calculated adsorption energy is specific to a particular binding site of a particular surface of a particular catalyst. Thus, the relationship between DFT-calculated adsorption energies and a catalyst's activity is not simple.

For example, in cases of lower adsorbate coverage on the catalyst surface, adsorbates tend to adsorb to stronger-binding sites before weaker-binding sites. In cases of higher adsorbate coverage, adsorption energies are difficult to calculate, so it is not uncommon to assume low adsorbate coverage.<sup>29–31</sup> We make this assumption here. Specifically, we assume that the adsorption energy of a surface is equivalent to the strongest adsorption energy seen across all sites on that surface. We then feed this adsorption energy into an activity volcano to predict the activity of the surface. Note that the activity volcano we used here<sup>32</sup> used changes in free energy ( $\Delta G$ ) to predict activity, while the datasets we used here contained changes in enthalpy ( $\Delta H$ ). We assumed that free energy corrections remained constant at 0.5 eV such that  $\Delta G = \Delta H + 0.50$  eV, as explained in our previous work.<sup>14</sup>

Given the activities of the surfaces of a catalyst, the next step is to estimate the activity of the entire catalyst. One way to do this would be to perform a weighted average of the surface activities, where higher weights are given to surfaces that are more stable. For the sake of simplicity, we instead propose a uniform average. We recognize that future work may involve investigating more sophisticated averaging methods. Such methods may be especially important because the exponential nature of the catalyst activity often leads to orders-of-magnitude shifts in activity between different surfaces. Thus, the estimates of bulk activity may be sensitive to the weights used for high-activity surfaces.

Concretely, suppose we have *n* catalyst candidates  $\{x_i\}_{i=1}^n$ , where each candidate  $x_i$  has *m* surfaces  $\{u_{i,j}\}_{j=1}^m$ , and the surface  $u_{i,j}$ 



FIG. 1. Illustration of Myopic Multiscale Sampling (MMS). Given a database of DFT-calculated adsorption energies (a), we train a ML model to predict adsorption energies (b). Then, we use those adsorption energies to estimate activities of catalyst surfaces (c), which we then use to estimate the activities of the bulk catalysts (d); choose which catalyst to sample next (e); choose which surface on the catalyst to sample (f); choose which site on the surface to sample (g); and perform DFT of that site to add to the database (h). This procedure is repeated continuously with the goal of classifying all catalysts as either "relatively active" or "relatively inactive."

has  $\ell$  sites  $\{s_{i,j,k}\}_{k=1}^{\ell}$ . For a given site  $s_{i,j,k}$ , denote its free energy of adsorption by  $\Delta G(s_{i,j,k})$ , and for a given surface  $u_{i,j}$ , denote its catalytic activity by  $\alpha(u_{i,j})$ . Likewise, for a given catalyst material candidate  $x_i$ , denote its average catalytic activity by  $\alpha(x_i) = \frac{1}{m} \sum_{j=1}^{m} \alpha(u_{i,j})$ . Suppose we have a predictive uncertainty estimate for the adsorption energy  $\Delta G(s_{i,j,k})$  of a site, represented by a normal distribution with mean  $\mu_{i,j,k}$  and variance  $\sigma_{i,j,k}^2$ . We can then perform simulation-based uncertainty quantification of catalyst activity by using the multiscale modeling process we described above to propagate uncertainties from sites' adsorption energies. Specifically, for each material candidate  $x_i$ , we generate H samples of its catalytic activity,  $\{\tilde{\alpha}_i^h\}_{h=1}^H$ , by simulating from the following generative process:

$$\begin{aligned} &\text{for } j = 1, \dots, m, \ k = 1, \dots, \ell, \\ & \{ \tilde{\Delta} G_{i,j,k}^h \}_{h=1}^{H} \stackrel{\text{iid}}{\sim} \mathcal{N} \Big( \mu_{i,j,k}, \sigma_{i,j,k}^2 \Big), \\ &\text{for } h = 1, \dots, H, \ j = 1, \dots, m, \\ & \tilde{\alpha}_{i,j}^h = \begin{cases} \exp(M_1 \tilde{\Delta} G_{i,j,1:\ell}^h + B_1) & \text{if } \tilde{\Delta} G_{i,j,1:\ell}^h \geq t^* \\ \exp(M_2 \tilde{\Delta} G_{i,j,1:\ell}^h + B_2) & \text{otherwise,} \end{cases}$$
(1)   
 &\text{for } h = 1, \dots, H, \\ & \tilde{\alpha}\_i^h = \frac{1}{m} \sum\_{i=1}^m \tilde{\alpha}\_{i,j}^h, \end{aligned}

where  $t^*$  is the optimal absorption energy for a given volcano relationship and  $M_1$ ,  $M_2$ ,  $B_1$ , and  $B_2$  are the linear coefficients associated with the two sides of the log-scaled volcano relationship of a given chemistry. Furthermore, we use the notation  $\tilde{\Delta} G_{i,i,1;\ell}^{h}$  to denote the set  $\{\tilde{\Delta}G_{i,j,k}^h\}_{k=1}^\ell$ . In other words, we first generate H samples from the adsorption energy distributions for each adsorption site. Then, we pair each of the H samples from one adsorption site to other samples from sites on the same surface. Within each set of pairings, we find the minimum/strongest/most negative adsorption energy and treat that as one sample for the adsorption energy of the surface. This leads to H samples for the adsorption energy of a surface. Next, we map these H adsorption energies for each surface through a literature-sourced volcano relationship<sup>32</sup> to generate H samples of activity for each surface. Due to the noncontinuous and opaque nature of the volcano relationship, we opt to simplify this mapping process by approximating the volcano as the junction of two linear approximations [Eq. (1)]. Last, we pair each of the H activity samples from each surface with corresponding activity samples from other surfaces within the same bulk and then average them to generate a sample for bulk activity. This leads to H samples for bulk activity, which can then be composed into a distribution of bulk activity. Figure 2 illustrates this process of using our multiscale modeling method to estimate catalyst

J. Chem. Phys. **154**, 124118 (2021); doi: 10.1063/5.0044989 Published under license by AIP Publishing



FIG. 2. Multiscale modeling strategy for estimating the activity of a catalyst. For each adsorption site, we obtain a machine-learned estimate of its adsorption energy along with uncertainty; the resulting distributions are shown as dotted lines. Then, we aggregate the energy distributions for all sites within each surface by performing a paired minimum operation across sites, to estimate the low-coverage adsorption energy for each surface. The distribution of this value is shown as a solid line. Next, we transform the energy distributions for all surfaces into activities using a Sabatier relationship. Finally, we average all the surface activities to obtain an estimate of overall catalyst activity.

activity from DFT-calculated adsorption energies, including uncertainty quantification.

Each catalyst material candidate  $x \in \mathcal{X}$  has some true catalytic activity level  $\alpha(x)$ . Our goal will be to determine the top p-% of catalyst material candidates in terms of their activity levels, which we denote  $\mathcal{X}_p = \{x \in \mathcal{X} : r(\alpha(x)) \ge \lfloor \frac{pn}{100} \rfloor\}$ , where  $r : \mathbb{R}_+ \to \{1, \ldots, n\}$ is a function mapping the activity level  $\alpha(x)$  to an index denoting its rank (from highest to lowest activity). Given a specified p, if a candidate material is in this set, i.e.,  $x_i \in \mathcal{X}_p$ , then we say that its associated binary value  $y_i = 1$ , or say  $y_i = 0$  otherwise. In simpler terms, we want to find the top p-% most active catalysts. For this paper, we choose p = 10% arbitrarily. Any catalyst that falls within the top 10% in terms of activity will be labeled active, and anything below the top 10% will be labeled inactive.

We can, therefore, frame our goal as determining the associated binary value  $y_i$  for each catalyst material candidate  $x_i \in \mathcal{X} = \{x_i\}_{i=1}^n$ . Suppose we have formed point estimates for each of the binary values, written as  $\{\hat{y}_i\}_{i=1}^n$ . To assess the quality of this set of estimates with respect to the set of true candidate values, we focus on the  $F_1$ score—a popular metric for classification accuracy, defined as

$$F_{1} = 2 \times \frac{\text{precision} \times \text{recall}}{\text{precision} + \text{recall}}$$
$$= \frac{2 \sum_{i=1}^{n} y_{i} \hat{y}_{i}}{2 \sum_{i=1}^{n} y_{i} \hat{y}_{i} + \sum_{i=1}^{n} (1 - y_{i}) \hat{y}_{i} + \sum_{i=1}^{n} y_{i} (1 - \hat{y}_{i})}.$$
(2)

Given a set of ground-truth values  $\{y_i\}_{i=1}^n$ , we are able to compute the  $F_1$  score for a chosen set of value estimates  $\{\hat{y}_i\}_{i=1}^n$ .

However, in practice, we will typically not have access to these ground-truth values and, thus, cannot compute this score in an online procedure. For use in online experiments, we will take advantage of a metric that yields an estimate of the change in  $F_1$  score. This metric is computable using only our model of the activity of each catalyst, without requiring access to ground-truth values  $\{y_i\}_{i=1}^n$ , and can be used to assess and compare the convergence of our methods. Furthermore, it can be used to provide an early stopping method for our active procedures. We will show experimentally in Sec. III that this metric shows a strong correlation with the  $F_1$  score.

#### **B.** Sampling strategy

The goal of MMS is to discover catalysts that are likely to be experimentally active. Optimization of catalytic activity is not the main priority because we assume that experimentalists have the ability to investigate multiple candidate catalysts and that unforeseen experimental issues may make some of these redundant. Instead, a greater focus is given on identification of a large number of candidates rather than finding "the most active" candidate. That is why the core sequential learning algorithm we use in MMS is active classification.<sup>25</sup> <sup>26</sup> To be specific, we use Level Set Estimation (LSE) to identify catalysts for DFT sampling. After identifying catalysts for DFT sampling, we then need to choose which surface of the catalyst to sample; here, we use techniques from active regression to choose a given surface to query (note that, here, regression is used simply as a subroutine within our ultimate task of active classification). Once a surface is chosen, we then attempt to find the strongest binding site on that surface by using active optimization of the adsorption energies. Thus, we combine three different sequential learning strategies across three different length scales to decide which sitebased DFT calculation will help us classify active vs inactive catalysts (Fig. 3).

We first describe the initial step of our sampling strategy, which consists of selecting a catalyst material candidate from our candidate set  $\mathcal{X} = \{x_i\}_{i=1}^n$ . Note that our high-level goal is binary classification in which we want to efficiently produce accurate estimates  $\{\hat{y}_i\}_{i=1}^n$  of the binary value for each material candidate. Based on our definition of  $y_i = \mathbb{I}[x_i \in \mathcal{X}_p]$ , this problem can be equivalently viewed as the task of LSE in which we aim to efficiently produce an accurate estimate of the superlevel set  $\mathcal{X}_p = \{x \in \mathcal{X} : r(\alpha(x)) \ge \lfloor \frac{pn}{100} \rfloor\}$ . There has been a body of work on developing acquisition functions for choosing candidates to query in the task of LSE.<sup>33,34</sup> In particular, we focus on the *probability of incorrect classification* acquisition function,<sup>35</sup> defined for an  $x_i \in \mathcal{X}$  as

 $\varphi(x_i) = \min(p, 1-p),$ 

where

$$p = \Pr\left(r(\alpha(x)) \ge \left\lfloor \frac{pn}{100} \right\rfloor\right)$$
$$\approx \frac{1}{H} \sum_{h=1}^{H} \mathbb{I}\left[r(\tilde{\alpha}_{i}^{h}) \ge \left\lfloor \frac{pn}{100} \right\rfloor\right].$$
(3)

Empirical probability  $\alpha(x)$  in top p-%



FIG. 3. Myopic Multiscale Sampling (MMS) overview. At the highest level, we choose a catalyst to query using level-set estimation—specifically, we use the probability of incorrect classification as our acquisition function. At the middle level, we choose a surface of the catalyst using uncertainty sampling. At the lowest level, we choose a site on the surface using Bayesian optimization to find the lowest energy site.

Thus, to select a subsequent catalyst candidate, we compute  $\varphi(x_i)$  for each  $x_i \in \mathcal{X}$  and return the maximizer  $x^* = \arg \max_{x_i \in \mathcal{X}} \varphi(x_i)$ . In simpler terms, we choose the catalyst that we are most likely to classify incorrectly. Note how this implies that we do *not* query catalysts that we are confident are active, which is different from active

optimization methods. This provides a more exploratory method i.e., focuses on bulks that are near the classification boundary rather than an exploitative one—i.e., focusing on bulks with highest activity, as would be done for an optimization task. Given our classification objective designed for early-stage computational discovery and screening, the former method is more appropriate.

ARTICLE

The selection of a catalyst candidate  $x_i$  depends on its estimated catalytic activity, which we model as an average of the catalytic activities across the surfaces of the candidate, i.e.,  $\alpha(x_i) = \frac{1}{m} \sum_{j=1}^{m} \alpha(u_{i,j})$ . Though we select a candidate based on its ability to help improve our estimate of the superlevel set  $\mathcal{X}_p$ , once selected, we then wish to most efficiently improve our estimate of this candidate's catalytic activity. Our goal at this stage is, therefore, to most efficiently learn the catalytic activities for each surface of that candidate. This can be viewed as an active regression task, where we aim to sample a surface that will most reduce the uncertainty of our surface activity estimates. To select a surface, we use an *uncertainty sampling for regression* acquisition function from the active learning literature,<sup>36</sup> defined as

$$\varphi(u_{i,j}) = \operatorname{Var}[\operatorname{Pr}(\alpha(u_{i,j}))]$$
$$\approx \frac{1}{H-1} \sum_{h=1}^{H} \left( \tilde{\alpha}_{i,j}^{h} - \frac{1}{H} \sum_{h'=1}^{H} \tilde{\alpha}_{i,j}^{h'} \right)^{2}, \qquad (4)$$

which selects a surface  $u_i^*$  of the material candidate  $x_i$  that has the greatest variance. In simpler terms, we choose the surface of a catalyst that has the most uncertainty because we suspect that this choice is most likely to reduce our uncertainty estimate of catalyst activity.

The catalytic activity of a given surface  $\alpha(u_{i,j})$  is a function of the adsorption energies of the sites on this surface, according to the relationship  $\alpha(u_{i,j}) = \exp(-|M\tilde{\Delta}G_{i,j,1:\ell} + B|)$  from Eq. (1), where  $\tilde{\Delta}G_{i,j,1:\ell}$  is the set of adsorption energies over all sites on the surface. Therefore, given a selected surface  $u_{i,j}$ , we wish to determine efficiently the site on this surface with minimum adsorption energy. This can be viewed as an optimization task. We, therefore, use the *expected improvement* acquisition function from Bayesian optimization,<sup>37</sup> defined as

$$\varphi(s_{i,j,k}) = \mathbb{E}\Big[ (\Delta G(s_{i,j,k}) \le \Delta G^*) \mathbb{1}\Big[ \Delta G(s_{i,j,k}) - \Delta G^* \Big] \Big] \\ \approx \Phi\bigg( \frac{\Delta G^* - \tilde{\mu}_{i,j,k}}{\tilde{\sigma}_{i,j,k}} \bigg) \varphi\bigg( \frac{\Delta G^* - \tilde{\mu}_{i,j,k}}{\tilde{\sigma}_{i,j,k}} \bigg) \big( \Delta G^* - \tilde{\mu}_{i,j,k} \big), \quad (5)$$

where  $\tilde{\mu} = \frac{1}{H} \sum_{h=1}^{H} \Delta \tilde{G}_{i,j,k}^{h}$  is the expected adsorption energy,  $\tilde{\sigma} = \sqrt{\frac{1}{H-1} \sum_{h=1}^{H} \left(\Delta \tilde{G}_{i,j,k}^{h} - \tilde{\mu}\right)^{2}}$  is its standard deviation,  $\Phi$  is the cumulative density function (CDF) of a standard normal distribution,  $\phi$  is the probability density function (PDF) of a standard normal distribution, and  $\Delta G^{*}$  is the minimum observed adsorption energy. This selects a site  $s_{i,j}^{*}$  that is expected to most reduce the site adsorption energy relative to the current minimum observed energy and allows for efficient estimation of the minimum energy site on the surface  $u_{i,j}$ . In simpler terms, we choose the site on a surface that is most likely to help us identify the strongest/lowest energy binding site on the surface.

#### C. Active learning stopping criteria

Assessing convergence of an active algorithm is useful for enabling early stopping, which can save resources. Measures of convergence can also provide diagnostics in online use settings. To quantify convergence, we use the *predicted change in F1 score*  $(\Delta \hat{F})$ .<sup>38</sup> Intuitively speaking, this rule says to stop an active learning procedure when  $\Delta \hat{F}$  drops and stays below a predefined threshold  $\epsilon$  for  $\kappa$  consecutive windows. We refer to  $\kappa$  as the stopping horizon, i.e.,

*Stop* if  $\Delta \hat{F} < \epsilon$  for  $\kappa$  consecutive iterations

Continue otherwise.

In our setting,  $\Delta \hat{F}$  is defined to be

$$\hat{\Delta}F = 1 - \frac{2a}{2a+b+c},\tag{6}$$

where *a* is the number of bulks for which the models at iterations *i* and *i* + 1 both yield a positive label, *b* is the number of bulks for which the model at iteration *i* yields a positive label, while that at iteration *i* + 1 yields a negative label, and *c* is the number of bulks for which the model at iteration *i* yields a negative label, while that at iteration *i* + 1 yields a positive label. Each of *a*, *b*, and *c* is computed over the previous *k* iterations. This measure provides an estimate of the change in accuracy at each iteration, and it allows us to control how conservatively (or aggressively) we stop early via an interpretable parameter *c*. We show results of this measure along-side our *F*1 score in Sec. III. Note that Altschuler and Bloodgood<sup>38</sup> recommend using a *stop set* of unlabeled points over which to calculate  $\Delta \hat{F}$ . Here, we use the entire search space of catalysts *in lieu* of a stop set that was representative of the search space.

#### D. Management of data queries

Implementation of MMS also involves definition of several hyper-parameters. For example, most surrogate models require training data before making predictions to feed the sampling method. This means that we needed to seed MMS with initial training data. We chose to create the initial training data by randomly sampling 1000 adsorption energies from the search space. We used random sampling for simplicity, and we sampled 1000 adsorption energies because that was the minimum amount of data on which the Convolution-Fed Gaussian Process (CFGP) (described below in more detail) could train on and maintain numerical stability.

Another consideration for MMS is the *batch size* and how to handle queries in tandem. Normal sequential learning assumes that we can make one query at a time. However, in applications such as ours, it may be possible to make multiple queries in parallel i.e., we can perform multiple DFT calculations at a time. There are several methods for handling queries in parallel; we chose to use a type of look-ahead sampling.<sup>39</sup> With look-ahead sampling, we began by choosing the first point to sample using the standard acquisition strategy. Then, while that point was still "being queried," we assumed that the first point was queried successfully and set the "observed" value equal to our predicted value. In other words, we pretend that we sampled the first data point and that our prediction of it was perfect. This allowed us to then recalculate our acquisition values to choose a second point. This process of "looking ahead" one point at a time was continued until a predetermined number of points were selected for querying—i.e., the batch size. Here, we chose a batch size of 200 points because that was roughly the number of DFT calculations that we could perform in a day during our previous high-throughput DFT studies.<sup>14</sup> Note that we did not retrain the surrogate models within each batch of 200 points; we only re-calculated acquisition values between each sample within each batch. We skipped re-training of surrogate models within each batch to reduce the amount of model training time required to perform this study. Although this may have reduced the effectiveness of the look-ahead method, we found the increased algorithm speed to be worthwhile.

#### E. Estimating performance through simulation

We aim to experimentally assess the performance of MMS and compare it with that of a variety of baseline methods without incurring the high cost of repeated DFT calculations. To do this, we simulate each procedure using a database of pre-determined adsorption energies. Specifically, suppose we have chosen a set of *n* catalyst material candidates  $\{x_i\}_{i=1}^n$  of interest. For each candidate  $x_i$ , we already have all the adsorption energies  $\Delta G(s_{i,j,k})$  for the full set of sites across the full set of surfaces on  $x_i$ . We can then run our procedures in a relatively fast manner, where we can quickly query the database at each iteration of a given method rather than running DFT. Similar offline-data discovery procedures have been pursued by previous work in optimization and active learning, where expensive evaluations have been collected offline and used for rapid online evaluation.<sup>40–42</sup>

One notable baseline method is *random search*, which at each iteration samples sites to carry out DFT calculations uniformly at random from the full set of sites over all catalyst material candidates. We provide simulation results using random search as a benchmark to compare MMS against.

## 1. Surrogate models used

Our objective in this paper is to assess the performance of MMS. The performance of MMS is likely to depend on the surrogate model used to predict adsorption energies from atomic structures. We assume that surrogate models with high predictive accuracy and calibrated uncertainty estimates<sup>43</sup> will outperform models with low accuracy and uncalibrated uncertainty estimates, but we are unsure of the magnitude of this difference. We, therefore, propose to pair at least two different models with MMS: a "perfect" model and an "ignorant" model.

We define the "perfect" model, hereby referred to as the "prime" model, as a model that returns the true adsorption energy of whatever data point is queried. This perfect prediction ensures high model accuracy. When asked for a standard deviation in the prediction, the prime model will return a sample from a  $\chi^2$  distribution whose mean is 0.1 electron volts (eV). This uncertainty ensures a sharp and calibrated<sup>43,44</sup> measure of uncertainty. We do not use standard deviation of zero because (1) it causes numerical issues during multiscale modeling and (2) any model in practice should not be returning standard deviations of zero.

We define the "ignorant" model, hereby referred to as the "null" model, as a model that returns the optimal adsorption energy no matter what is queried. This constant prediction ensures a relatively low model accuracy. When asked for a standard deviation in the prediction, the null model will return 1 eV. This uncertainty ensures a relatively dull and uncalibrated measure of uncertainty.

Last, we also choose to use a third, most practical model: CFGP.<sup>44</sup> The CFGP is a Gaussian process regressor whose features are the output of the final convolutional layer in a trained graph convolutional neural network. This model is our best current estimate of both an accurate and a calibrated model that could be used in practice. Thus, we have three models: null, CFGP, and prime, which are intended to give quantitative estimates of the minimal, medial, and maximal performances of MMS, respectively.

### 2. Search spaces used

Previous studies have shown that different materials discovery problems have varying difficulties.<sup>18</sup> Searching for a needle in a hay stack is generally more difficult than searching for a leaf on a branch. Thus, any simulation we do depends on the search space we use. To obtain a range of potential MMS performances, we perform simulations using two different datasets. Both datasets comprise thousands of atomic structures that represent CO adsorbing onto various catalyst surfaces, as well as corresponding adsorption energies. We then use Sabatier relationships from the literature to transform the adsorption energies into estimates of activity.<sup>32</sup>

We defined our first search space by synthesizing it randomly. We did so by retrieving a database of enumerated adsorption sites from the Generalized Adsorption Simulator for Python (GASpy).<sup>14,45</sup> These sites composed all the unique sites on all surfaces with Miller indices between -2 and 2 across over 10 000 different bulk crystal structures. We then randomly selected 200 of the bulk crystals along with all of the resulting surfaces and sites, yielding over 390 000 adsorption sites. Then, for each bulk crystal, we randomly sampled its "bulk mean adsorption energy" from a unit normal distribution. For each surface within each crystal, we randomly sampled its "surface mean adsorption energy" from a normal distribution whose mean was centered at the corresponding bulk mean and whose standard deviation was set to 0.3 eV. Then, for each site within each surface, we randomly sampled its adsorption energy from a normal distribution whose mean was centered at the corresponding surface mean and whose standard deviation was set to 0.1 eV. Thus, the adsorption energies were correlated within each bulk, and they were also correlated within each surface. We used a standard deviation of 1.0 eV for the catalysts because it was similar to the ~1 eV variability in energy shifts we saw in our GASpy dataset. We used a standard deviation of 0.1 eV for the sites because DFT precision does not often exceed 0.1 eV. We used a standard deviation of 0.3 eV for the surfaces because it was between the 1.0 and 0.1 eV standard deviations of the catalysts and sites, respectively.

We defined our second search space by retrieving our database of ~19 000 DFT-calculated CO adsorption energies calculated by GASpy, hereafter referred to as the GASpy dataset. The sites in this database were chosen using previous iterations of our sequential learning methods,<sup>14</sup> and they, therefore, have bias in the locations at which they were sampled. Specifically, the sites in this database were chosen based on the likelihood that their adsorption energies were close to the optimal value of -0.67 eV.<sup>14,32</sup>

There are several advantages of using the synthesized dataset over the real GASpy dataset, and vice versa. The synthesized dataset contains pseudo-random adsorption energies that are difficult for the CFGP to predict, thereby hindering its performance unfairly. Therefore, we should not and did not use the CFGP with the synthesized dataset; we used it with the GASpy dataset only. On the other hand, the number of surfaces per bulk and the number of sites per surface in the GASpy dataset were relatively sparse compared to those in the synthesized dataset. This can result in catalysts that require relatively few site queries to sample fully, which reduces the number of queries necessary to classify a catalyst. This reduction in the number of required queries per catalyst could artificially improve the observed performance of MMS.

## **III. RESULTS**

At the beginning of the simulations, the multiscale models made their catalyst class predictions (i.e., active or inactive) using the adsorption energy predictions and uncertainties of the models. As the simulations progressed and adsorption energies were queried, the models' predictions of each queried energy were replaced with the "true" value of the query and the corresponding uncertainty was collapsed to 0 eV. This was done to mimic a realistic use case where we would not use model predictions when we had the "real" DFT data instead. It follows that as the simulations progressed and nearly all points were queried, most models performed similarly because they all had comparable amounts of "true" data to use in the multiscale model. Conversely, the number of "predictions" at the end of the simulations decreased, and so differences between methods also decreased as simulations progressed.

#### A. Performance on synthesized data

This behavior is seen in Fig. 4(a), which shows how the F1 score changes at each point in the simulation of the synthesized dataset. Here, we see that the simulations using the prime model began with an F1 score of ~0.6 that increased to 1 over time. On the other hand, simulations using the null model began with an F1 score closer to 0 or 0.2 before gradually increasing to 1. This shows that more accurate surrogate models for adsorption energies led to more accurate multiscale models, even initially. Note also that the rate at which the F1 score improved was better when using MMS than when using random sampling, especially when using the null model. These data may suggest that the rate of improvement is governed by the acquisition strategy, while the initial performance is governed by the model.

Figure 4(b) shows how the  $\Delta \hat{F}$  changes at each point in the simulation of the synthesized dataset. The simulations using random search generally yielded higher  $\Delta \hat{F}$  values. This indicates slower convergence, which is consistent with the slower F1 score increase seen in the random search curves in Fig. 4(a). Note also how the  $\Delta \hat{F}$  values for the MMS-prime simulation decreased at around 500 batches, which is the number of batches it took for the F1 score to reach ~1. Last, we note that the  $\Delta \hat{F}$  values for the MMS-null simulation were often zero. This is because the null model was a "stiff" learner that did not result in any multiscale modeling changes unless a lowcoverage adsorption site was found. This shows that slow-learning models may result in relatively low  $\Delta \hat{F}$  values, which may necessitate longer stopping horizons (i.e.,  $\kappa$ ) to offset this behavior. In other words, worse models may need longer horizons before stopping the discovery to mitigate the chances of missing important information.

# The Journal of Chemical Physics



FIG. 4. Performance and convergence results for the simulations on the synthesized dataset. (a) *F1* score of the multiscale model during simulation of the synthesized data. For clarity of visualization, we plotted the rolling average of the *F1* score using a window of 20 batches. (b)  $\Delta \hat{F}$  of the multiscale model during simulation of the synthesized data. For clarity of visualization, we plotted the rolling average of  $\Delta \hat{F}$  using a window of 40 batches (excluding the MMS null line, where no averaging was done). RS represents "random search," while MMS represents myopic multiscale sampling.

These simulations provided us with an estimate of the improvement in active classification that we may get from using MMS. With the synthesized dataset, we saw that the MMS-with-null case achieved an F1 score of ~0.6 after ~250 batches (or 50 000 queries). This was over seven times faster than the random-sample-withnull case, which achieved an F1 score of ~0.6 after ~1800 batches (or 360 000 queries). When using the prime model, MMS was able to achieve an F1 score of ~0.75 in 200 batches, while the random search achieved the same performance in ~1200 batches, or six times slower.

We remind readers that we allowed each method to supplant model predictions with "true" values from their respective queries. This is why all methods converged on identically perfect performance near the end of the simulations, where each method had near-perfect information to use.

## B. Performance on DFT data

Figure 5 shows the *F*1 score and the  $\Delta \hat{F}$  of the multiscale model at each point in the simulation of the GASpy dataset. Interestingly, the system performance when using the CFGP was similar to the performance when using the null model, both of which were overshadowed by the relatively good performance when using the prime



FIG. 5. Performance and convergence results for the simulations on the GASpy dataset. (a) *F*1 score of the multiscale model during simulation of the GASpy dataset. (b)  $\Delta \hat{F}$  of the multiscale model during simulation of the GASpy dataset. RS represents "random search," while MMS represents myopic multiscale sampling.

model. This suggests that there is a large room for improvement for the CFGP model. Note also how the MMS strategy outperforms random sampling for this dataset as well.

These simulations provided us with a second estimate of the improvement in active classification that we may get from using MMS. With the GASpy dataset, we saw that the MMS-with-null case achieved an F1 score of ~0.8 after ~6 batches (or 1200 queries). This was over 16 times faster than the random-sample-with-null case, which achieved an F1 score of ~0.6 after ~80 batches (or 16000 queries). When using the prime model, both MMS and random search were able to achieve an F1 score of ~0.8 after only a single batch.

To validate our multiscale model, we used it to classify the 5% most active catalysts in the GASpy dataset. We found that copper alloys composed a third of this set of most active catalysts, which is consistent with the research community's findings that copper is one of the most active elements for reducing  $CO_2$ .<sup>46,47</sup> We also found aluminum, tin, and gallium alloys in this active set of alloys, which have also appeared in  $CO_2$  reduction literature.<sup>15,48,49</sup> Yttrium,

scandium, and silicon alloys also appear in this active set. We acknowledge that these results are based on analyses of a dataset that was built using more heuristic methods, and further investigation is needed to verify the integrity of these candidates. However, at minimum, these results are not inconsistent with current  $\rm CO_2$  reduction literature.

#### C. Recommended diagnostics

We note that the *F*1 scores illustrated in Figs. 4(a) and 5(a) cannot be calculated without knowing all the true classes, which are not possible to know during a real discovery process. We need metrics to monitor the behavior of both our discovery algorithms. We recommend monitoring the  $\Delta \hat{F}$  as well as the accuracy, calibration, and sharpness (i.e., the magnitude of the predicted uncertainties) of the surrogate model over time. Figure 6 shows an example of such diagnostic metrics over the course of our simulation that used MMS and CFGP on the GASpy dataset.



FIG. 6. Example of diagnostic plots that we recommend monitoring during an active discovery campaign: (a) predicted change in F1 score  $(\Delta \hat{F})$ ; (b) residuals between the real data and the surrogate model's predictions; (c) expected calibration error<sup>44</sup> of the surrogate model; (d) predicted uncertainties of the surrogate model in the form of the predicted standard deviation ( $\sigma$ ); (e) negative-log-likelihood of the surrogate model.44 These results were simulated by using the Myopic Multiscale Sampling (MMS) method with the Convolution-Fed Gaussian Process (CFGP) model on the GASpy dataset. For clarity of visualization, we plotted rolling averages of all values in this figure using a window of 100 queries (excluding the  $\Delta \hat{F}$  values, where no averaging was done).

 $\Delta \hat{F}$  estimates the amount of overall improvement in the discovery process. Sustained low values of  $\Delta \hat{F}$  are a necessary but not sufficient indicator of convergence. To improve our confidence in the predictive strength of  $\Delta \hat{F}$ , we can test one of its underlying assumptions: the multiscale model becomes progressively more accurate as it receives more data. This assumption is true when we replace surrogate model predictions with incoming DFT results, but it is not necessarily true for unqueried points. We can estimate the accuracy on unqueried points by calculating the residuals between the surrogate model and the incoming DFT results [Fig. 6(b)]. As each "batch" of queries is received, we compare the queried, true adsorption energies with the energies predicted by the surrogate model just before retraining-i.e., the predictions used to choose that batch. Any improvements in accuracy on these points show that the overall, multiscale model is improving over time and that the  $\Delta \hat{F}$  metric is an honest indicator of convergence. Figure 6(b) shows that model accuracy improves within the first ~10 batches (or 2000 adsorption energy queries) but plateaus afterward. This indicates that, after ten batches, improvements in overall classification accuracy came from receipt of additional DFT data rather than improvements in surrogate model predictions.

Prediction accuracy of adsorption energies is not the only indicator of improved model performance. If a surrogate model's accuracy does not change but its uncertainty predictions decrease/improve, then our confidence in the overall material classification may still improve. Of course, improvements in uncertainty must not be obtained at the expense of worse calibration. In other words, reductions in predicted uncertainties may also indicate improved model performance and better confidence in  $\Delta \hat{F}$ , but only if the expected calibration error<sup>44</sup> does not increase. In our illustrative example, Fig. 6(c) shows the predicted uncertainty predictions do not decrease over the course of the discovery process. Note that all uncertainty and calibration estimates for each batch should be calculated using the surrogate model predictions used to choose that batch, just as was done for the residuals.

Last, we also recommend monitoring the negative-loglikelihood<sup>44</sup> of the surrogate model for each incoming batch. This metric incorporates model accuracy, calibration, and sharpness into a single metric. Lower values of negative-log-likelihood indicate better model performance. Figure 6(e) shows that this metric improves until ~2000 queries, after which it stagnates. This is consistent with the improvement in accuracy until 2000 queries and subsequent stagnation of all performance metrics thereafter.

## **IV. CONCLUSIONS**

Here, we created a multi-scale modeling method for combining atomic-scale DFT results with surrogate/ML models to create actionable plans for experimentalists—i.e., a classification of catalysts as "worthy of experimental study" or "not worthy." We then coupled this modeling method with a Myopic Multiscale Sampling (MMS) strategy to perform automated catalyst discovery via active classification. We tested this strategy on two hypothetical datasets using three different surrogate models, giving us an estimate on the range of performances we might see in the future. In some cases, the results show up to a 16-fold reduction in the number of DFT queries compared to random sampling. The degree of speed-up depends on the quality of the ML model used, the homogeneity of the search space, and the hyperparameters used to define convergence of the active classification. Speed-up estimates on more realistic use cases show a more conservative seven-fold reduction in the number of DFT queries. Finally, we provide a set of recommended diagnostic metrics to use during active classification (Fig. 6):  $\Delta \hat{F}$  and the ML model's residuals, uncertainty estimates, and calibration.

Our results elucidated a number of qualitative behaviors of active classification. First, we observed that higher-quality ML models yielded better initial performance of the classification process. Conversely, we observed that higher-quality sampling strategies yielded better rates of improvement over time. We also observed that our latest ML model (CFGP) yielded performance closer to a naive, ignorant model than to a perfect, omniscient model. This suggests that there is a relatively large amount of potential improvement left in the ML modeling space. Next, we observed that better sampling strategies (as quantified by the F1 score) led to lower rates of change in classes (as quantified by  $\Delta \hat{F}$ ), suggesting that  $\Delta \hat{F}$ may be an indicator of sampling strategy performance. Conversely, we observed that slow-learning ML models may also reduce  $\Delta \hat{F}$ . This phenomenon could be counteracted by using more conservative convergence criteria. All these details were observed in specific and synthetic use cases though. The behaviors seen here may not be observed in situations where search spaces and/or ML models differ.

We encourage readers to focus on the main goals of this work: (1) converting atomic-scale simulations and ML models into actionable decisions for experimentalists and (2) relaxing the active discovery process from an optimization/regression problem to a classification problem. The ability to convert computational results into experimental recommendations helps us serve the research community better. Simultaneously, relaxing the discovery process to a classification problem helps us prioritize exploration rather than exploitation, which is more appropriate for early-stage discovery projects.

We also recognize several future directions that may stem from this research. Future work might include incorporation of DFTcalculated surface stability by performing weighted averaging of surface activities when calculating bulk activities. Future work may also include cost-weighted sampling such that less computationally intensive calculations are chosen more frequently than more intensive ones, which may improve discovery rates in real-time. Perhaps most importantly, future work should incorporate some ability to feed experimental data and information to computational sampling strategies—e.g., multi-fidelity modeling.

# AUTHORS' CONTRIBUTIONS

Kevin Tran and Willie Neiswanger contributed equally to this work.

#### ACKNOWLEDGMENTS

This research used resources of the National Energy Research Scientific Computing Center, a DOE Office of Science User Facility supported by the Office of Science of the U.S. Department of Energy under Contract No. DE-AC02-05CH11231. K.B. acknowledges the U.S. Department of Energy, Office of Science, Office of Basic Energy Sciences, Data Science for Knowledge Discovery for Chemical and Materials Research program under Award No. DESC0020392. W.N. was supported by the U.S. Department of Energy Office of Science under Contract No. DE-AC02-76SF00515.

#### DATA AVAILABILITY

The data that support the findings of this study are openly available at https://github.com/ulissigroup/catalyst-acquisitions/.

#### REFERENCES

<sup>1</sup>A. Jain, S. P. Ong, G. Hautier, W. Chen, W. D. Richards, S. Dacek, S. Cholia, D. Gunter, D. Skinner, G. Ceder, and K. A. Persson, "The materials project: A materials genome approach to accelerating materials innovation," APL Mater. 1, 011002 (2013).

<sup>2</sup>S. Curtarolo, G. L. W. Hart, M. B. Nardelli, N. Mingo, S. Sanvito, and O. Levy, "The high-throughput highway to computational materials design," Nat. Mater. 12, 191–201 (2013).

<sup>3</sup>J. E. Saal, S. Kirklin, M. Aykol, B. Meredig, and C. Wolverton, "Materials design and discovery with high-throughput density functional theory: The open quantum materials database (OQMD)," JOM **65**, 1501–1509 (2013).

<sup>4</sup>S. P. Ong, W. D. Richards, A. Jain, G. Hautier, M. Kocher, S. Cholia, D. Gunter, V. L. Chevrier, K. A. Persson, and G. Ceder, "Python Materials Genomics (pymatgen): A robust, open-source python library for materials analysis," Comput. Mater. Sci. 68, 314–319 (2013).

<sup>5</sup>H. M. Berman, T. Battistuz, T. N. Bhat, W. F. Bluhm, P. E. Bourne, K. Burkhardt, Z. Feng, G. L. Gilliland, L. Iype, S. Jain, P. Fagan, J. Marvin, D. Padilla, V. Ravichandran, B. Schneider, N. Thanki, H. Weissig, J. D. Westbrook, and C. Zardecki, "The protein data bank," Nucl. Acids Res. 28, 235–242 (2000).

<sup>6</sup>T. Meyer, M. D'Abramo, A. Hospital, M. Rueda, C. Ferrer-Costa, A. Pérez, O. Carrillo, J. Camps, C. Fenollosa, D. Repchevsky, J. L. Gelpí, and M. Orozco, "MoDEL (molecular dynamics extended library): A database of atomistic molecular dynamics trajectories," Structure 18, 1399–1409 (2010).

<sup>7</sup>P. Schlexer Lamoureux, K. T. Winther, J. A. Garrido Torres, V. Streibel, M. Zhao, M. Bajdich, F. Abild-Pedersen, and T. Bligaard, "Machine learning for computational heterogeneous catalysis," ChemCatChem **11**, 3581–3601 (2019).

<sup>8</sup>G. R. Schleder, A. C. M. Padilha, C. M. Acosta, M. Costa, and A. Fazzio, "From DFT to machine learning: Recent approaches to materials science—-a review," J. Phys.: Mater. **2**, 032001 (2019).

<sup>9</sup>A. J. Medford, M. R. Kunz, S. M. Ewing, T. Borders, and R. Fushimi, "Extracting knowledge from data through catalysis informatics," ACS Catal. **8**, 7403–7429 (2018).

<sup>10</sup>Y. G. Chung, J. Camp, M. Haranczyk, B. J. Sikora, W. Bury, V. Krungleviciute, T. Yildirim, O. K. Farha, D. S. Sholl, and R. Q. Snurr, "Computation-ready, experimental metal-organic frameworks: A tool to enable high-throughput screening of nanoporous crystals," Chem. Mater. **26**, 6185–6192 (2014).

<sup>11</sup>E. I. Ioannidis, T. Z. H. Gani, and H. J. Kulik, "molSimplify: A toolkit for automating discovery in inorganic chemistry," J. Comput. Chem. **37**, 2106–2117 (2016).

<sup>12</sup>S. Chakraborty, W. Xie, N. Mathews, M. Sherburne, R. Ahuja, M. Asta, and S. G. Mhaisalkar, "Rational design: A high-throughput computational screening and experimental validation methodology for lead-free and emergent hybrid perovskites," ACS Energy Lett. 2, 837–845 (2017).
 <sup>13</sup>M. L. Green, C. L. Choi, J. R. Hattrick-Simpers, A. M. Joshi, I. Takeuchi, S. C.

<sup>15</sup>M. L. Green, C. L. Choi, J. R. Hattrick-Simpers, A. M. Joshi, I. Takeuchi, S. C. Barron, E. Campo, T. Chiang, S. Empedocles, J. M. Gregoire, A. G. Kusne, J. Martin, A. Mehta, K. Persson, Z. Trautt, J. Van Duren, and A. Zakutayev, "Fulfilling the promise of the materials genome initiative with high-throughput experimental methodologies," Appl. Phys. Rev. 4, 011105 (2017).

 $^{14}$  K. Tran and Z. W. Ulissi, "Active learning across intermetallics to guide discovery of electrocatalysts for CO<sub>2</sub> reduction and H<sub>2</sub> evolution," Nat. Catal. 1, 696–703 (2018).

<sup>15</sup> M. Zhong, K. Tran, Y. Min, C. Wang, Z. Wang, C.-T. Dinh, P. De Luna, Z. Yu, A. S. Rasouli, P. Brodersen, S. Sun, O. Voznyy, C.-S. Tan, M. Askerka, F. Che, M. Liu, A. Seifitokaldani, Y. Pang, S.-C. Lo, A. Ip, Z. Ulissi, and E. H. Sargent, "Accelerated discovery of CO<sub>2</sub> electrocatalysts using active machine learning," Nature **581**, 178–183 (2020).

<sup>16</sup>T. Lookman, P. V. Balachandran, D. Xue, and R. Yuan, "Active learning in materials science with emphasis on adaptive sampling using uncertainties for targeted design," npj Comput. Mater. 5, 21 (2019).

<sup>17</sup>J. Schmidt, M. R. G. Marques, S. Botti, and M. A. L. Marques, "Recent advances and applications of machine learning in solid-state materials science," npj Comput. Mater. 5, 83 (2019).

<sup>18</sup>Y. Kim, E. Kim, E. Antono, B. Meredig, and J. Ling, "Machine-learned metrics for predicting the likelihood of success in materials discovery," npj Comput. Mater. 6, 131 (2020); arXiv:1911.11201.

<sup>19</sup>B. Rohr, H. S. Stein, D. Guevarra, Y. Wang, J. A. Haber, M. Aykol, S. K. Suram, and J. M. Gregoire, "Benchmarking the acceleration of materials discovery by sequential learning," Chem. Sci. 11, 2696–2706 (2020).

<sup>20</sup>B. Settles, Synthesis Lectures on Artificial Intelligence and Machine Learning, edited by R. J. Brachman, W. W. Cohen, and T. G. Dietterich (Morgan & Claypool, 2012), p. 100.

<sup>21</sup> P. I. Frazier, "A tutorial on bayesian optimization," arXiv:1807.02811v1 (2018).
 <sup>22</sup> S. C. H. Hoi, D. Sahoo, J. Lu, and P. Zhao, "Online learning: A comprehensive survey," arXiv:1802.02871 (2018).

<sup>23</sup>A. Khorshidi and A. A. Peterson, "Amp: A modular approach to machine learning in atomistic simulations," Comput. Phys. Commun. 207, 310–324 (2016).

<sup>24</sup>J. Vandermause, S. B. Torrisi, S. Batzner, Y. Xie, L. Sun, A. M. Kolpak, and B. Kozinsky, "On-the-fly active learning of interpretable Bayesian force fields for atomistic rare events," npj Comput. Mater. 6, 20 (2020).

<sup>25</sup>Y. Ma, D. J. Sutherland, R. Garnett, and J. Schneider, "Active pointillistic pattern search," in 18th International Conference on Artificial Intelligence and Statistics (AISTATS) (PMLR, 2015), Vol. 38, pp. 672–680.

<sup>26</sup>A. Zanette, J. Zhang, and M. J. Kochenderfer, "Robust super-level set estimation using Gaussian processes," Lect. Notes Comput. Sci. **11052 LNAI**, 276–291 (2019); arXiv:1811.09977.

<sup>27</sup>R. A. Flores, C. Paolucci, K. T. Winther, A. Jain, J. A. G. Torres, M. Aykol, J. Montoya, J. K. Nørskov, M. Bajdich, and T. Bligaard, "Active learning accelerated discovery of stable iridium oxide polymorphs for the oxygen evolution reaction," Chem. Mater. **32**, 5854–5863 (2020).

<sup>28</sup>Z. W. Seh, J. Kibsgaard, C. F. Dickens, I. Chorkendorff, J. K. Nørskov, and T. F. Jaramillo, "Combining theory and experiment in electrocatalysis: Insights into materials design," Science 355, eaad4998 (2017).

<sup>29</sup>J. K. Nørskov, F. Studt, F. Abild-Pedersen, and T. Bligaard, Angewandte Chemie International Edition (John Wiley & Sons, 2015).

<sup>30</sup> J. K. Nørskov, T. Bligaard, A. Logadottir, J. R. Kitchin, J. G. Chen, S. Pandelov, and U. Stimming, "Trends in the exchange current for hydrogen evolution," J. Electrochem. Soc. **152**, J23 (2005).

<sup>31</sup>E. M. Lopato, E. A. Eikey, Z. C. Simon, S. Back, K. Tran, J. Lewis, J. F. Kowalewski, S. Yazdi, J. R. Kitchin, Z. W. Ulissi, J. E. Millstone, and S. Bernhard, "Parallelized screening of characterized and DFT-modelled bimetallic colloidal Co-catalysts for photocatalytic hydrogen evolution," ACS Catal. 10, 4244–4252 (2020).

<sup>32</sup>X. Liu, J. Xiao, H. Peng, X. Hong, K. Chan, and J. K. Nørskov, "Understanding trends in electrochemical carbon dioxide reduction rates," Nat. Commun. 8, 15438 (2017).

<sup>33</sup>A. Gotovos, N. Casati, G. Hitz, and A. Krause, "Active learning for level set estimation," in *Proceedings of the Twenty-Third International Joint Conference on Artificial Intelligence* (AAAI Press, 2013), pp. 1344–1350.

<sup>34</sup> K. Kandasamy, W. Neiswanger, R. Zhang, A. Krishnamurthy, J. Schneider, and B. Poczos, "Myopic posterior sampling for adaptive goal oriented design of experiments," in *International Conference on Machine Learning* (PMLR, 2019), pp. 3222–3232.
 <sup>35</sup> B. Bryan, R. C. Nichol, C. R. Genovese, J. Schneider, C. J. Miller, and L.

<sup>55</sup>B. Bryan, R. C. Nichol, C. R. Genovese, J. Schneider, C. J. Miller, and L. Wasserman, "Active learning for identifying function threshold boundaries," in *Advances in Neural Information Processing Systems* (2005), Vol. 18, pp. 163–170.

pp. 163–170. <sup>36</sup>B. Settles, "Active learning literature survey," Technical Report 1648, University of Wisconsin-Madison, Department of Computer Sciences, 2009.

<sup>37</sup>J. Močkus, "On Bayesian methods for seeking the extremum," in *Optimization Techniques IFIP Technical Conference* (Springer, 1975), pp. 400–404.

J. Chem. Phys. **154**, 124118 (2021); doi: 10.1063/5.0044989 Published under license by AIP Publishing <sup>38</sup>M. Altschuler and M. Bloodgood, "Stopping active learning based on predicted change of F measure for text classification," in *Proceedings of 13th IEEE International Conference on Semantic Computing, ICSC 2019* (IEEE, 2019); arXiv:1901.09118, Vol. 47–54.

<sup>39</sup>T. Desautels, A. Krause, and J. Burdick, "Parallelizing exploration-exploitation tradeoffs with Gaussian process bandit optimization," J. Mach. Learn. Res. 15, 4053–4103 (2014); arXiv:1206.6402.

<sup>40</sup>I. Char, Y. Chung, W. Neiswanger, K. Kandasamy, A. O. Nelson, M. D. Boyer, E. Kolemen, and J. Schneider, "Offline contextual Bayesian optimization for nuclear fusion," in *33rd Conference on Neural Information Processing Systems* (*NeurIPS*) (2020); arXiv:2001.01793, pp. 1–12.

<sup>41</sup>C. Ying, A. Klein, E. Real, E. Christiansen, K. Murphy, and F. Hutter, "NAS-BENCH-101: Towards reproducible neural architecture search," in *36th International Conference on Machine Learning* (PMLR, 2019); arXiv:1902.09635.

<sup>42</sup>C. White, W. Neiswanger, and Y. Savani, "Bananas: Bayesian optimization with neural architectures for neural architecture search," arXiv:1910.11858 (2020). <sup>43</sup>V. Kuleshov, N. Fenner, and S. Ermon, "Accurate uncertainties for deep learning using calibrated regression," in *35th International Conference on Machine Learning* (PMLR, Stockholm, Sweden, 2018); arXiv:1807.00263.

<sup>44</sup> K. Tran, W. Neiswanger, J. Yoon, Q. Zhang, E. Xing, and Z. W. Ulissi, "Methods for comparing uncertainty quantifications for material property predictions," Mach. Learn.: Sci. Technol. 1, 025006 (2020).

<sup>45</sup>K. Tran, A. Palizhati, S. Back, and Z. W. Ulissi, "Dynamic workflows for routine materials discovery in surface science," J. Chem. Inf. Model. 58, 2392–2400 (2018).
<sup>46</sup>J. Qiao, Y. Liu, F. Hong, and J. Zhang, Chem. Soc. Rev. 43, 631–675 (2014).

<sup>47</sup>B. Khezri, A. C. Fisher, and M. Pumera, "CO<sub>2</sub> reduction: The quest for electrocatalytic materials," J. Mater. Chem. A **5**, 8230–8246 (2017).

<sup>48</sup>X. Zheng, Y. Ji, J. Tang, J. Wang, B. Liu, H.-G. Steinrück, K. Lim, Y. Li, M. F. Toney, K. Chan, and Y. Cui, "Theory-guided Sn/Cu alloying for efficient CO<sub>2</sub> electroreduction at low overpotentials," Nat. Catal. 2, 55–61 (2019).

 $^{49}$  D. A. Torelli, S. A. Francis, J. C. Crompton, A. Javier, J. R. Thompson, B. S. Brunschwig, M. P. Soriaga, and N. S. Lewis, "Nickel-gallium-catalyzed electro-chemical reduction of CO<sub>2</sub> to highly reduced products at low overpotentials," ACS Catal. 6, 2100–2104 (2016).