

Bayesian optimization of single-atom alloys and other bimetallics:

Efficient screening for alkane transformations, CO₂ reduction, and hydrogen evolution

Gbolade O. Kayode,^a Avery F. Hill,^a and Matthew M. Montemore.^a

^a Department of Chemical and Biomolecular Engineering, Tulane University, New Orleans, LA 70118, USA

Abstract

Single-atom alloys form an important class of material that has shown great potential in maximizing the use of rare and expensive metals in catalysis due to their high catalytic performance, robustness, tunability and unique structure. Single-atom alloys present particular challenges for screening as they can deviate from some traditional catalyst design frameworks. While machine learning (ML) can be quite useful in accelerating catalyst design, most traditional ML methods require relatively large datasets and/or high-level, expensive featurization. Additionally, most of these ML methods are incapable of handling multiple objectives and constraints over the intended search space. In this work, we leverage Bayesian optimization (BO) to guide our search for high-performing catalysts. We show that our BO workflow can be initialized with as few as 2 to 8 data points, and often identifies the optimal single-atom alloy surface in just a few iterations. Our workflow was used to efficiently search across multiple adsorption systems and datasets and significantly outperformed a random search method, using simple, off-the-shelf features. For applications, we used BO to identify potential high-performing catalysts for alkane transformations, CO₂ reduction, and hydrogen evolution. Our BO workflow identified Hf₁Cu for alkane transformations; Y₁Au, Y₁Cu and Y₁Ag for CO₂ reduction; and an Ag-Ir bimetallic alloy for hydrogen evolution. Simple stability tests indicate all three single-atom alloys for CO₂ reduction are stable and most likely synthesizable. The workflow developed here can also be used for experiments or high-level theory calculations as well as other classes of materials.

1 Introduction

In the past decade, single-atom alloys (SAAs) have emerged as a new frontier in catalysis and materials science more generally.^{1–7} This is largely due to their unique structure, high atom utilization efficiency and distinctive interfacial sites.^{1,5,8} SAAs have also been shown to break the constraint of scaling relations, which can result in high catalytic performance.^{9,10} They have also been shown to be highly robust, stable and tunable, allowing for high activity and selectivity.^{1,3–5,8,11} SAAs are created when small amounts of a metal (usually an active metal) is atomically dispersed into the surface layer of another metal (usually a coinage metal) resulting in a unique, well-defined, single-site structure. The dispersed metal can serve as an active site for catalytic reactions, allowing for precise control over the reaction pathways and enhancing catalytic performance.^{1,12}

As with traditional alloys, machine learning (ML) models have also been used in the rational design of SAAs. This is due to the large compositional space of alloys as well as the cost associated with experiments and high-level theory calculations. For instance, ML has been used to screen SAAs for acetylene semi-hydrogenation,¹³ and identified 5 SAAs with high predicted selectivity and activity. In another study, ML was used to predict oxygen and carbon adsorption energies to screen for ethanol dehydrogenation performance.¹⁴ Other studies show ML being used to predict SAA properties such as stability and segregation energies.^{15–18} Overall, these studies demonstrate the utility of ML in the rapid screening of SAAs for different applications. However, for SAAs (or other SACs), it would be desirable to have a ML method that efficiently screens smaller design spaces, in contrast to most ML methods that require a relatively large dataset. For example, for a given host metal there are less than 30 possible transition-metal dopants; this is a relatively small search space. Also, the uncertainty of most ML methods is not well quantified,

which can result in inaccurate predictions of high performance or missing out on truly high-performing candidates. Furthermore, many ML approaches require high-level featurization,^{19–21} which sometimes requires additional experiments or quantum calculations. There is therefore a need for efficient search techniques that can comb through both small and large spaces effectively and without significant feature engineering, thus minimizing the amount of computational time and researcher effort. This is useful not only in accelerating standard DFT-based screening but is crucial in cases where generating data is more difficult or expensive, such as experiments or high-level calculations such as the random-phase approximation.

Bayesian optimization (BO) is one such technique. BO is particularly useful because it efficiently samples a given search space, eliminating the need for a priori generation of a large number of experiments or DFT calculations. The central idea behind BO is the use of a prior distribution in acquiring a posterior distribution and subsequently selecting the next point(s) of evaluation from the posterior distribution by use of an acquisition function. By intelligently choosing points of evaluation, BO can often achieve its objective in fewer iterations in comparison to other optimization techniques. This makes BO ideal in the evaluation of expensive objective functions. BO has excelled in many domains^{22–24} and is one of the underlying technologies behind most self-driving labs.^{25–27} One major advantage of BO over traditional ML algorithms for catalyst discovery is its capacity to handle multiple objectives and constraints over the search space. Additionally, because BO is a form of active learning, it does not necessarily require a high form of featurization to perform effectively. Some studies have used BO for material discoveries in computational chemistry; however, only few of these studies use BO for adaptive experimental design. Also, there are relatively few studies that involve search spaces with discrete quantities²⁸ as opposed to continuous functions. Other Bayesian techniques such as the Bayesian information

criterion²⁹ and Bayesian model averaging³⁰ have been employed by previous studies;³¹ however, these techniques are used to aid in creating models on an existing dataset rather than for recommending new data points (i.e., adaptive experimental design), which is the main purpose of BO.

In this work, we leveraged BO for adaptive experimental design in the theoretical discovery of SAAs. Our results show that our BO workflow can efficiently search for materials with target adsorption energies through multiple types of adsorption systems and datasets. Our BO workflow outperformed a random search method, confirming its efficiency for material discovery. We also show that our BO framework works well with simple, off-the-shelf features, eliminating the need for high-level featurization. We applied our BO workflow to efficiently identify potential high-performing catalysts for alkane transformations, CO₂ reduction, and hydrogen evolution. Within a few iterations, our BO workflow identified Hf₁Cu as a promising candidate for alkane transformations. For CO₂ reduction, Y₁Au, Y₁Cu and Y₁Ag were identified, and a subsequent simple stability test for these three candidates suggests that they are stable and are likely to be synthesizable. To demonstrate our BO workflow is quite efficient even with a relatively large design space, we employed it in the search for promising bimetallics for the hydrogen evolution reaction (HER). Our BO workflow identified an Ag-Ir bimetallic alloy as a promising candidate with high predicted performance. In general, we showcase the efficiency of BO for adaptive experimental design in materials discovery, particularly for SAAs.

2 Datasets

The adsorption energy datasets used in this work were taken from previous studies, with the partial exception of the CO adsorption energy dataset where the initial data was taken from previous work but new calculations were performed as part of the BO search. The first dataset

(methane-intermediate dataset)³² contains adsorption energies of methane decomposition intermediates (C, H, CH, CH₂, CH₃) on Cu-based SAAs. This dataset was generated from DFT calculations using the VASP software with the Perdew-Burke-Ernzerhof (PBE) functional^{33,34} and the project-augmented wave (PAW) potential.^{35,36} The SAAs were constructed using a 4-layer Cu(111) 3 x 3 supercell, with the bottom two layers fixed. One atom on the surface layer of this slab was then substituted with 47 metals ranging from Li (atomic number 3) to Bi (atomic number 81). The adsorption energies of C, H, CH, CH₂, and CH₃ were then computed on each surface by taking the difference between the total energy of the slab-adsorbate system and the sum of the total energies of the bare slab and the geometrically relaxed adsorbate. Broadly, the methane-intermediate dataset represents adsorption energies from the same host metal, but different adsorbates.

In contrast to the previous dataset, the second dataset (H dataset)³⁷ represents adsorption energies from different host metals, but with the same adsorbate. This dataset contains H adsorption energies on Cu-based, Ag-based and Au-based SAAs. The adsorption energies in this dataset were also calculated using the VASP software, but in this case the Perdew-Burke-Ernzerhof (PBE) functional together with the Tkatchenko–Scheffler method³⁸ for van der Waals interactions were used. Also, all calculations were spin-polarized, and the adsorption energies were calculated relative to gas-phase H₂. In total, there were 21 Cu-based, 22 Ag-based and 21 Au-based adsorption energies. We also used a bimetallic dataset³¹ consisting of 2934 H adsorption energies adsorbed on different sites. Top, bridge, fcc-hollow, and hcp-hollow sites were all present for each bimetallic in this dataset.

The last dataset (CO dataset)³⁹ contains a total of 22 CO adsorption energies on 10 Cu-based, 8 Ag-based and 4 Au-based SAAs. The computational setup used here was identical to that

used for the H dataset. This setup was also used for the additional DFT calculations we performed in this work for CO adsorption energies and SAA stabilities.

3 Results and Discussion

3.1 Workflow for Bayesian Optimization




Figure 1: BO workflow for optimizing the reactivity of single atom alloys.

Our BO workflow for SAA catalyst discovery is shown in Figure 1. The objective of this workflow is to identify a SAA that has an adsorption energy that falls near a target value. First, we chose the search space and collected the initial set of DFT calculations, which was then used to construct the initial surrogate model. In this case, we used the standard Gaussian process regressor (GPR) as our surrogate model. GPR is a probabilistic model used to predict the distribution of the objective function across the search space. It is commonly used for BO due to its ability to generate uncertainties over predicted values. In constructing the GPR model, we used a new feature combination composed of elemental data taken from literature. Due to its convenience and effectiveness, this approach of using existing elemental data is very common among ML-based

screening studies.^{14,40} Our feature set included the group number (G)⁴⁰, period number (P)⁴⁰, coupling matrix element (V_{ad}^2)^{9,41,42}, and oxygen adsorption energy on a pure surface composed of the dopant metal (O_{ads}). We found that including O_{ads} (taken from literature⁴¹) somewhat improved the efficiency of our searches; however, we also found that BO is still reasonably efficient when O_{ads} is not used (we discuss this further in the next section). Next, we computed the acquisition function using the distribution and uncertainties generated from the GPR model. For our acquisition function, we used the expected improvement, which is widely recognized as the standard choice. The acquisition function was then used to generate the next recommendation. Specifically, the workflow recommends the SAA with the maximum expected improvement. Afterwards, we performed the DFT calculations for the recommended SAA (or, in most cases, pulled them from the existing database) and checked if the stopping criteria were met. If not, the training set was augmented with the recommendation and steps 3-6 were repeated until the stopping criteria was met. We have two stopping criteria: The loop ends (1) if the adsorption energy of the recommendation falls within 2% of the target adsorption energy; or (2) if after 14 DFT calculations, the target adsorption energy does not fall within the uncertainty region of the unexplored SAAs. 14 was chosen as an arbitrary threshold. We also employ contextual improvement⁴³ to control the exploration vs. exploitation trade-off. This allows the BO framework to dynamically choose what the exploration factor should be after each iteration based on prior information. This approach has been shown in previous work⁴³ to be quite efficient.

To concretely demonstrate our workflow, we considered an example where the goal was to identify a dopant for a Cu-based SAA with an overall carbon adsorption energy of -6.23 eV, $\pm 2\%$. First, we selected the initial set of calculations by choosing 8 Cu-based SAAs (see Supporting information) from the same-host dataset, as shown in Figure 2a (green). This initial set

was chosen by hand to sample across the periodic table; as we demonstrate below the average efficiency is not greatly sensitive to the attributes of the initial set. The search space consists of an additional 26 Cu-based SAAs (see Supporting information). Next, we constructed the GPR model using the initial set of calculations and their features. We then predicted the distribution and uncertainties over the search space using the constructed model. This distribution was then used to compute the acquisition function, and this function's maximum was identified to generate the first recommendation, which was Mn (Figure 2a, blue). Mn_1Cu gave a carbon adsorption energy of -7.17 eV (see Figure 2b). At this point, neither of the two stopping criteria is met, so the loop continues. The second recommendation was Os (Figure 2a, blue), with a carbon adsorption energy of -7.50 eV, and the third recommendation was Sc (Figure 2a, blue and red) with a carbon adsorption energy of -6.17 eV. Because Sc_1Cu has a carbon adsorption energy within 2% of the target energy (-6.23 eV), the loop ends. Thus, it took just 3 additional DFT calculations to achieve our objective in this case. If no valid solution existed in the search space, the second stopping criterion would be invoked after 14 DFT calculations (arbitrary threshold) if none of the untested SAAs had uncertainties that overlap with the target space.




Figure 2: Example BO test: The goal is to identify a Cu-based SAA that gives an overall carbon adsorption energy of $-6.23 \text{ eV} \pm 2\%$. (a) Periodic table showing the initial set of calculations and each recommendation in order. (b) Carbon adsorption energy of each recommendation. Sc_1Cu with a carbon adsorption of -6.17 eV was identified after 3 DFT calculations.

3.2 BO Tests

The example discussed in the previous section showcases the utility of our BO workflow in guiding searches by providing intelligent recommendations. To more broadly demonstrate our workflow, we performed a series of search campaigns across multiple datasets involving randomly generated target energies.

First, we considered 5 search campaigns for each of the adsorbates (C, H, CH, CH_2 , and CH_3) in the methane-intermediate dataset (same host metal, various adsorbates). The goal of these campaigns was to identify Cu-based SAAs that fall within $\pm 2\%$ of their respective target energies. The 8 initial calculations and 26 additional candidates in the search space were the same as in

Figure 2. Figure 3a shows the average number of DFT calculations and the standard deviations across all 5 campaigns for each of the adsorbates. An average of number of 3.2, 4, 3, 4.2, and 6.4 additional DFT calculations were required for C, H, CH, CH₂, and CH₃ respectively. This set of campaigns have thus shown our BO workflow to be quite efficient, identifying target SAAs with a relatively small number of DFT calculations.

Next, we evaluated how well our BO workflow performed in searching across various dopants for each host metal. Similar to the previous set of search campaigns, we considered 5 campaigns for each host metal (Cu, Ag, Au) in the H dataset (same adsorbate, various host metals). The goal was again to identify SAAs that fall within $\pm 2\%$ of their respective target energies. In this case, the initial set of calculations contained 6 SAAs for both Ag- and Au-based SAAs and 5 for the Cu-based SAAs (see Supporting information). The search space for this set of campaigns had 16 Ag-based SAAs, 15 Au-based SAAs and 16 Cu-based SAAs (see Supporting information). The average additional DFT calculations (Figure 3b) needed for Ag, Au and Cu were 4, 4.4 and 3.4 respectively.

Further, we considered a combination of Au-based, Ag-based and Cu-based SAAs from the previous campaign (Figure 3b). We combined all the initial calculations from the previous Au, Ag and Cu search campaigns together to serve as our initial set for this campaign, for a total of 17 SAAs. Afterwards, we combined the search space for all 3 hosts from the previous search campaigns together, totaling 47 SAAs for these campaigns. The targets were also randomly chosen, and a dummy feature was added to the feature set to represent the 3 different host metals. In this case, an average of 4.4 DFT calculations were required to achieve the targets (Figure 3c, green).

To further test the efficiency of our BO workflow, we compared it to a random search method. The random search method arbitrarily chooses a recommendation from the search space until the target energy is acquired. We used this search method on the same exact search campaigns as the mixed-host dataset (Figure 3c). The random search method required an average of 19.6 DFT calculations to reach the target energies (Figure 3c, yellow), more than 4 times more than BO. This further illustrates the efficiency of our BO workflow.

We also examined the influence of the O_{ads} feature (the O adsorption energy on the dopant as a pure metal, taken from literature⁴¹) on the search efficiency of the same exact campaigns from Figure 3c. To achieve this, we considered two feature sets: (1) our previously used feature set with O_{ads} , G , P , V_{ad}^2 and a dummy variable (to represent host metals), and (2) the same feature set except with O_{ads} removed. Based on Figure 3d, including O_{ads} allows the BO to converge with an average of 4.4 DFT calculations, and without it an average of 10.2 DFT calculations are needed. This shows that O_{ads} improves the efficiency, but without it the scheme is still significantly more efficient than the random search.

Lastly, to further test how sensitive the BO efficiency is to the feature set, we ran the same 5-fold search campaign for CH but with a completely different feature set. This time our feature set was made up of the heat of formation of the dopant metal, dipole polarizability of the dopant metal, Pauling electronegativity of the dopant metal, and fusion heat of the dopant metal. These features were arbitrarily chosen from the Mendelev python database.⁴⁴ In this case, our BO workflow took an average of 4.8 additional DFT calculations to converge. This is slightly higher than the 3 additional calculations needed when using our original feature set (see Figure S1). Thus, we demonstrate our BO workflow is still effective even with a different, arbitrarily chosen feature set, although the efficiency does depend somewhat on the feature set.




Figure 3: Series of BO search campaigns. In each case, the average number of DFT calculations (not including the initial seed set) needed to achieve the target adsorption energy is shown, with the standard deviations shown by error bars. (a) 5-fold BO search campaign across all adsorbates in the methane-intermediate dataset, which has the same metal host (Cu) but various adsorbates. (b) 5-fold BO search campaign across each metal host in the H adsorption dataset. (c) BO vs random search across a combination of all three host metals from the H adsorption dataset. (d) 5-fold BO search campaign performed with O_{ads} and without O_{ads} included in the feature set. The BO workflow is still reasonably efficient without O_{ads}.

3.3 The Effect of the Initial set on BO Efficiency

Overall, we have shown the effectiveness and robustness of our BO framework in handling different types of search campaigns for target adsorption energies. However, because the search efficiency of our BO campaigns depends on the initial set of DFT calculations, it is useful to understand which properties of this initial set of DFT calculations affect the BO campaign efficiency. Broadly, this provides useful insight into how to design an efficient BO campaign.

Thus, we carried out a set of BO tests to examine how the size of the initial set of DFT calculations, the dispersion of initial dopants on the periodic table, the number of initial dopants in the same group, and the number of initial dopants in the same period all affect the efficiency of our BO searches. We considered the same search campaign as Figure 2, where the goal was to identify a Cu-based SAA that had an adsorption energy within 2% of -6.23 eV. If a BO search could not find its target, the search was associated with the maximum number of DFT calculations, as defined by stopping criterion 2. Overall, as described below, collections of different initial sets were used to run a total of 1200 BO searches for the same search campaign in Figure 2.

First, because BO is an iterative process, we wanted to see if the size of the initial set of calculations had a significant impact on the search efficiency. To test this, we generated 100 initial sets of calculations for each size ranging from 2 to 6 using random sampling without replacement. Figure 4a shows the average number of DFT calculations in the BO search campaign in blue, while the green bars show the inherent cost of running more initial DFT calculations to obtain the initial set. Interestingly, the sample size appears to have no effect on BO campaign efficiency (mean additional calculations is 5.3, 5.3, 5.7, 5.3, 5.0 for 2 – 6 initial calculations), which indicates that using smaller initial set sizes is likely to decrease the overall computational cost.

Next, we wanted to examine if dispersion across the periodic table would lead to an increased BO efficiency, as this may make the initial GPR model more broadly accurate across the search space. Using the previously generated initial sets of size 4, we obtained a measure of seed dispersion by calculating the Euclidean distance between dopant metals on the periodic table (using the samples' group number and period number to create a 2D coordinate system) and summing the distance between each pair of dopant metals within the initial set (see Figure 4b).

There is no apparent relationship between initial set dispersion on the periodic table and the BO campaign efficiency.

Furthermore, we examined the effect of having various period and group repetitions in the initial set of calculations (i.e., multiple dopants in the same period or group). This is largely because periodicity has been shown to be a good descriptor of adsorption energies on SAAs. We generated 100 initial sets of a fixed size (4) with 1 to 4 dopants in the same period and 1 to 3 dopants in the same group, then ran BO searches for all 700 of them and averaged the results. From Figure 4c, we see that the BO campaign efficiency does not depend on period repetitions. In contrast, Figure 4d shows a roughly proportional relationship between group repetition and BO campaign efficiency. This suggests that we could minimize computational cost by choosing initial sets with more dopants in the same group.




Figure 4: Series of BO search campaigns to ascertain the effect of (a) size of the initial set of calculations, (b) the dispersion of elements in the initial set of calculations, (c) the number of repetitions of elements along the same period on the periodic table (d) the number of repetitions of elements along the same group in the periodic table. Only group repetitions have a notable effect.

3.4 Applications

Due to the demonstrated high catalytic performance of SAAs for a variety of reactions, there continue to be intensive searches for promising SAAs for a variety of chemical processes. In this section, we applied our BO workflow to search for potentially effective catalysts for three important catalytic processes using simple design principles.

3.4.1 Alkane Transformations

The conversion of alkanes into more valuable products is a very important industrial process for both technological and environmental purposes.⁴⁵ This process is a vital building block for the production of plastics, synthetic rubber, synthesis gas, and a wide variety of useful chemicals.^{10,32,45–47} For example, methane steam reforming involves the catalytic conversion of methane to synthesis gas, which is a valuable raw material for methanol synthesis and Fischer-Tropsch synthesis.^{48,49} The ideal catalyst for alkane conversion reactions will activate the alkane (which can be measured via strong CH₃ adsorption) while simultaneously permitting the rapid formation and desorption of products (which can be measured via weak C adsorption) thereby inhibiting coking.^{10,50} However, the search for the optimal catalyst via the traditional approach is often difficult due to the linear scaling relations that often exist between C and CH₃ adsorption energies.

Here, we used BO to search for a promising Cu-based SAA catalyst for alkane conversion using the design principle previously stated (strong CH₃ adsorption but weak C adsorption). The goal was to identify a SAA that performed similarly or better than Ti₁Cu, which was found in

previous work¹⁰ to be a promising candidate for methane steam reforming. We used C and CH₃ adsorption energies from the methane-intermediate dataset. We maintained the same initial set of materials and feature set previously used for Figure 3a. For our target, we used a CH₃ adsorption energy of -2.6 eV and a C adsorption energy of -4.8 eV (see Figure 5). In this case, we are not aiming to achieve this particular target; instead, BO attempting to move towards this target will discover cases with strong CH₃ adsorption but weak C adsorption. The stopping criterion was that the loop ends if the recommendation has an equal or stronger CH₃ adsorption energy than Ti₁Cu (-1.97 eV), as well as an equal or weaker C adsorption energy than Ti₁Cu (-6.38 eV). The first recommendation, Ni₁Cu, does not satisfy our stopping criterion, but the second, Hf₁Cu, does satisfy the criterion, which ends the loop. Thus, in just two additional DFT calculations our BO workflow effectively optimizes through a space of 2D adsorption energies, making BO very useful in searching for materials that do not obey scaling relations and thus are challenging to design. Data from previous work indicates that Hf₁Cu is stable against aggregation and therefore is likely to be synthesizable.³⁹ In general, while Hf₁Cu appears to be a promising candidate for activating alkanes while preventing coking (based on the design principle we used), it may or may not turn out to be an effective catalyst for a particular alkane conversion reaction for a variety of reasons. For example, we did not consider selectivity, possible poisoning by other intermediates or contaminants, etc. Thus, further DFT or experimental testing would be needed.




Figure 5: BO search campaign for MSR. The goal is to identify a Cu-host SAA with both weaker C adsorption and stronger CH₃ adsorption as compared to Ti₁Cu. BO proceeds through this 2D search space of C and CH₃ adsorption energies. Ti is shown in blue, while the first and second recommendations (Ni and Hf) are shown in green, and the rest of the unexplored SAAs are shown in grey.

3.4.2 CO₂ reduction.

In the pursuit of sustainability and carbon neutrality, efficient conversion of CO₂ to useful products has received intensive research attention. By lowering net greenhouse gas emissions and offering a renewable source of carbon-based chemicals, this process could potentially play a significant role in combating climate change. However, many challenges remain, such as the design of catalysts with high efficiency and selectivity toward desired products. Previous work⁵¹ has identified Cu as the best pure-metal electrocatalyst for this process.

Here, the goal is to identify a SAA that is predicted to outperform Cu for CO₂ electroreduction using our BO workflow. We leveraged a volcano plot from previous work⁵¹ that can estimate the catalytic performance for this process using CO adsorption energies. Using the CO adsorption energy on Cu as reference, we scaled the volcano plot to match adsorption energy data from our computational setup. After scaling, the apex of the volcano plot had a CO adsorption

energy of -0.78 eV, compared to pure Cu at -0.89 eV. We then defined our target region as -0.78 ± 0.1 eV. For the initial set of calculations, we leveraged 22 CO adsorption energies consisting of Ag-based, Au-based, and Cu-based SAAs (see Supporting information) from our previous work.³⁹ Our search space contained a total of 75 SAAs consisting of Ag-based, Au-based Cu-based SAAs, with dopants ranging from transition metals to post-transition metals. For our feature set, we used the same feature set as in Figure 3c and d. The stopping criteria was similar to our previous campaigns with the target region of -0.78 ± 0.1 eV. In this case, we performed new DFT calculations as recommended by BO, rather than drawing them from a database.

The first, second, third and final recommendation were Bi₁Au, W₁Au, Tl₁Cu and Y₁Au with CO adsorption energies -0.12 eV, -2.58 eV, -0.12 eV and -0.77 eV respectively (see Figure 6a). Y₁Au (-0.77 eV) is predicted to be more active than pure Cu (-0.89 eV) for CO₂ reduction. Next, we decided to replace O_{ads} with C_{ads} (the C adsorption energy on a pure surface of the dopant metal) in our feature set. Previous work had suggested both O_{ads} and C_{ads} might be good descriptors for CO adsorption.^{49,52,53} We repeated the same BO procedure and found Y₁Cu (-0.69 eV) and Y₁Ag (-0.84 eV) at the 3rd and 8th iteration respectively (see Figure 6b). These results suggest that Y₁Au, Y₁Cu, and Y₁Ag are more active than pure Cu and would likely be suitable for this process. However, it is important to acknowledge that the descriptor-based approach has certain limitations, such as focusing primarily on activity. Thus, further tests for selectivity would need to be performed. Two recent theoretical studies^{6,54} have also predicted that Y₁Cu is indeed a good catalyst for this process, further corroborating our findings. By requiring a small number of DFT calculations to identify SAAs, this further demonstrates the utility and efficiency of our BO workflow for material discovery.




Figure 6: (a) CO₂ reduction screening using O_{ads} as part of the feature set. The BO workflow efficiently identified Y_1Au as a promising candidate for this reaction. (b) CO₂ reduction screening using C_{ads} instead of O_{ads} . The BO workflow efficiently identified Y_1Cu and Y_1Ag as promising candidates. Y_1Au , Y_1Cu and Y_1Ag are predicted to be more active than pure Cu.

Lastly, we performed simple stability tests of our promising candidates (see SI for further description of the stability tests). The reaction energy of $2Y_1M_{35} \rightarrow Y_2M_{34} + M_{36}$ must be positive for Y_1M_{35} to be stable against aggregation. Here, Y_1M_{35} is the SAA and M represents Au, Ag or Cu. This stability test is quite common and has been used in previous studies^{39,55}. The reaction energies were 1.17 eV, 0.69 eV and 1.27 eV for Y_1Au , Y_1Cu , and Y_1Ag respectively. This means Y_1Au , Y_1Cu , and Y_1Ag are stable against aggregation. Next, we carried out tests to determine dopant segregation. The energy difference, $\Delta E^{CO} = E_{bulk}^{CO} - E_{surface}^{CO}$ must be positive for the SAA to be stable against segregation in the presence of CO. The energy differences were 0.27 eV, 1.06 eV and 0.23 eV for Y_1Au , Y_1Cu , and Y_1Ag respectively. This indicates all three are stable against dopant segregation to the subsurface. Overall, Y_1Au , Y_1Cu , and Y_1Ag are predicted to be stable and would most likely be synthesizable.

3.4.3 Hydrogen evolution.

To demonstrate that our BO workflow can also efficiently screen through a relatively large search space, we screened through an existing dataset of H adsorption energies for possible active sites for HER. Here, the goal is to identify a bimetallic alloy that is predicted to outperform Pt for HER using our BO workflow. We leveraged a volcano plot from previous work⁵⁶ to estimate the catalytic activity for this process using H adsorption energies. To account for differences in computational setup, we shifted the volcano plot from this previous work using Pt as a reference point. After shifting, the apex of the volcano plot had an H adsorption energy of -0.225 eV (relative to gas-phase H₂), compared to pure Pt at -0.31 eV. We then defined our target region as $-0.225 \pm 10\%$ eV. We utilized just 2% of the bimetallic dataset to initialize the first BO iteration, which corresponds to 59 H adsorption energies on bimetallics in different sites. Our search space contained a total of 2875 H adsorption energies on bimetallics in various sites. We used the same feature set as before except for V_{ad}^2 , which we substituted in place of the *s-d* coupling of the dopant metal. To account for possible heterogeneity within a site, we took the weighted average of each feature for all neighboring surface atoms. This has been shown to be a good characterization of the adsorbing atom's local environment.⁴² We also added dummy variables to represent each type of site (top, bridge, fcc hollow, hcp hollow). The stopping criteria was similar to our previous campaigns.

BO identified a promising candidate for HER in just four iterations. The first, second, third and final recommendation were Au₉Zr₃ on the top site, Hf₃Zr₉ on the fcc-hollow site, Os₆W₆ on the top site, and Ag₆Ir₆ on the top site (see Figure 7). Ag₆Ir₆ (-0.226 eV) is predicted to be more active than pure Pt (-0.31 eV) for HER, based on the volcano plot. Ag-Ir alloys have also been shown to be synthesizable from previous work.⁵⁷ Thus, it is a strong candidate for further investigation. Therefore, we have shown our BO workflow performs quite well on a relatively

large search space. This test also shows that our BO workflow can perform well on materials other than SAAs. Despite the good performance of our sequential BO workflow on this large search space, batch BO (where $k > 1$ recommendation is provided per iteration) may be more effective for relatively large search spaces in general.




Figure 7: HER screening with BO. The BO workflow efficiently identified Ag₆Ir₆ on the top site as a promising candidate for this reaction.

4 Discussion and Conclusion

In general, the utility of our BO workflow is not limited to SAAs: our BO workflow is also capable of handling different types of systems such as single-atom catalysts (SACs)⁹, dual-atom catalysts (DACs)⁵⁸, dual-atom alloys (DAAs)³⁹, and other classes of materials. Furthermore, our BO workflow can be utilized for experiments and high-level theory calculations. More specifically, our BO workflow is particularly effective for learning from an arduously collected dataset (e.g., data from literature spanning many years), and making subsequent searches much more efficient. We demonstrated this with the search campaign for CO₂ reduction, where a dataset from previous work was used to accelerate a new search. We showed our BO technique to be

effective in relatively small search spaces, as this is a neglected (but important) use-case for machine learning. However, it is also effective for larger search spaces, as demonstrated by our HER screening.

In summary, we developed a BO workflow to efficiently guide the search of materials, specifically SAAs. We ran a series of BO search campaigns across multiple adsorption systems and datasets, demonstrating the efficiency of our BO workflow in identifying potentially high-performing catalysts with a small number of DFT calculations/iterations. Our BO workflow significantly outperformed a random search method, demonstrating its efficacy in searching through a complex space. We also demonstrated that our BO workflow can be used with simple, off-the-shelf features and does not necessarily require a high-level featurization like many traditional ML methods for material screening. Furthermore, we applied our BO workflow to identify optimal catalysts for three technologically important processes: alkane transformations, CO₂ reduction, and HER. With just two additional DFT calculations/iterations our BO workflow identified Hf₁Cu as an ideal candidate for alkane transformations. For CO₂ reduction, after very few DFT calculations/iterations, our BO workflow identified Y₁Au, Y₁Cu and Y₁Ag as promising candidates for this chemical process. Recent studies also corroborate our findings, as Y₁Cu was recently predicted via DFT to be an effective catalyst for this process, further demonstrating the effectiveness of our workflow. Simple stability tests also indicate Y₁Au, Y₁Cu and Y₁Ag are stable (against aggregation and segregation) and would most likely be synthesizable. For HER, our BO workflow identified Ag-Ir bimetallic alloy within a few iterations. Overall, our BO workflow can be applied to different classes of materials, experiments, and high-level theory calculations, and would be particularly useful in cases where data generation is expensive.

Conflicts of interest

There are no conflicts of interest to declare.

Acknowledgements

We thankfully acknowledge support from the U.S. – Israel Center for Fossil Fuels, administered by the BIRD foundation, support from Tulane University, and support from the National Science Foundation through grant CHE-2154952. G.O.K gratefully acknowledges the Connolly Alexander Institute for Data Science for their award. Computational resources for this work were provided by Tulane University's Technological Services and the Louisiana Optical Network Infrastructure ([http:// www.loni.org](http://www.loni.org)).

References

- 1 R. T. Hannagan, G. Giannakakis, M. Flytzani-Stephanopoulos and E. C. H. Sykes, *Chem. Rev.*, 2020, **120**, 12044–12088.
- 2 T. Zhang, A. G. Walsh, J. Yu and P. Zhang, *Chem. Soc. Rev.*, 2021, **50**, 569–588.
- 3 Y. Da, R. Jiang, Z. Tian, X. Han, W. Chen and W. Hu, *SmartMat*, 2023, **4**, 1–18.
- 4 Z. Xu, Z. Ao, M. Yang and S. Wang, *J. Hazard. Mater.*, 2022, **424**, 127427.
- 5 M. T. Darby, M. Stamatakis, A. Michaelides and E. C. H. Sykes, *J. Phys. Chem. Lett.*, 2018, **9**, 5636–5646.
- 6 N. Zhang, Y. Si, X. Chen, X. Wang and J. Yao, *ACS Appl. Nano Mater.*, 2023, **6**, 2394–2402.
- 7 T. D. Spivey and A. Holewinski, *J. Am. Chem. Soc.*, 2021, **143**, 11897–11902.
- 8 J. Mao, J. Yin, J. Pei, D. Wang and Y. Li, *Nano Today*, 2020, **34**, 100917.
- 9 G. O. Kayode, S. Zhang and M. M. Montemore, in *Catalysis*, 2022, vol. 34, pp. 17–55.
- 10 C. F. Nwaokorie and M. M. Montemore, *J. Phys. Chem. C*, 2022, **126**, 3993–3999.
- 11 X. Zhi, Y. Jiao, Y. Zheng, A. Vasileff and S.-Z. Qiao, *Nano Energy*, 2020, **71**, 104601.
- 12 R. Réocreux and M. Stamatakis, *Acc. Chem. Res.*, 2023, **2022**, 30.
- 13 H. Feng, H. Ding, S. Wang, Y. Liang, Y. Deng, Y. Yang, M. Wei and X. Zhang, *ACS Appl. Mater. Interfaces*, 2022, **14**, 25288–25296.
- 14 A. Kumar, J. Iyer, F. Jalid, M. Ramteke, T. S. Khan and M. A. Haider, *ChemCatChem*, 2022, **14**, e202101481.
- 15 G. A. Sulley, J. Hamm and M. M. Montemore, *J. Phys. Energy*, 2023, **5**, 015002.

- 16 M. Salem, M. J. Cowan and G. Mpourmpakis, *ACS Omega*, 2022, **7**, 4471–4481.
- 17 K. K. Rao, Q. K. Do, K. Pham, D. Maiti and L. C. Grabow, *Top. Catal.*, 2020, **63**, 728–741.
- 18 Z. K. Han, D. Sarker, R. Ouyang, A. Mazheika, Y. Gao and S. V. Levchenko, *Nat. Commun.* 2021 121, 2021, **12**, 1–9.
- 19 R. García-Muelas and N. López, *Nat. Commun.*, 2019, **10**, 1–7.
- 20 M. Andersen and K. Reuter, *Acc. Chem. Res.*, 2021, **54**, 8.
- 21 M. M. Montemore and J. W. Medlin, *J. Phys. Chem. C*, 2014, **118**, 2666–2672.
- 22 B. Shahriari, K. Swersky, Z. Wang, R. P. Adams and N. de Freitas, *Proc. IEEE*, 2016, **104**, 148–175.
- 23 M. Diessner, J. O’Connor, A. Wynn, S. Laizet, Y. Guan, K. Wilson and R. D. Whalley, *Front. Appl. Math. Stat.*, 2022, **8**, 119.
- 24 Y.-F. Lim, C. K. Ng, U. S. Vaitesswar and K. Hippalgaonkar, *Adv. Intell. Syst.*, 2021, **3**, 2100101.
- 25 M. Abolhasani and E. Kumacheva, *Nat. Synth.* 2023, 2023, 1–10.
- 26 B. P. MacLeod, F. G. L. Parlane, C. C. Rupnow, K. E. Dettelbach, M. S. Elliott, T. D. Morrissey, T. H. Haley, O. Proskurin, M. B. Rooney, N. Taherimakhsoosi, D. J. Dvorak, H. N. Chiu, C. E. B. Waizenegger, K. Ocean, M. Mokhtari and C. P. Berlinguette, *Nat. Commun.*, 2022, **13**, 1–10.
- 27 F. Häse, L. M. Roch and A. Aspuru-Guzik, *Trends Chem.*, 2019, **1**, 282–291.
- 28 T. Kim, J. Ahn, N. Kim and S. Yun, *Adaptive Local Bayesian Optimization Over Multiple Discrete Variables*, 2020.
- 29 A. A. Neath and J. E. Cavanaugh, *WIREs Comput. Stat.*, 2012, **4**, 199–203.
- 30 C. T. Volinsky, A. E. Raftery, D. Madigan and J. A. Hoeting, *Stat. Sci.*, 1999, **14**, 382–417.
- 31 O. Mamun, K. T. Winther, J. R. Boes and T. Bligaard, *npj Comput. Mater.*, 2020, **6**, 177.
- 32 T. Toyao, K. Suzuki, S. Kikuchi, S. Takakusagi, K. Shimizu and I. Takigawa, *J. Phys. Chem. C*, 2018, **122**, 8315–8326.
- 33 J. P. Perdew, J. A. Chevary, S. H. Vosko, K. A. Jackson, M. R. Pederson, D. J. Singh and C. Fiolhais, *Phys. Rev. B*, 1992, **46**, 6671–6687.
- 34 J. Perdew, K. Burke and M. Ernzerhof, *Phys. Rev. Lett.*, 1996, **77**, 3865–3868.
- 35 P.E.Blöchl, *Phys. Rev. B*, 1994, **50**, 17953–17979.
- 36 G. Kresse and D. Joubert, *Phys. Rev. B*, 1999, **59**, 1758–1775.
- 37 A. P. Monasterial, C. A. Hinderks, S. Viriyavaree and M. M. Montemore, *J. Chem. Phys.*,

- 2020, **153**, 111102.
- 38 A. Tkatchenko and M. Scheffler, *Phys. Rev. Lett.*, 2009, **102**, 073005.
- 39 S. Zhang, E. C. H. Sykes and M. M. Montemore, *Chem. Sci.*, 2022, **13**, 14070–14079.
- 40 X. Li, R. Chiong and A. J. Page, *J. Phys. Chem. Lett.*, 2021, **12**, 5156–5162.
- 41 G. O. Kayode and M. M. Montemore, *J. Mater. Chem. A*, 2021, **9**, 22325–22333.
- 42 M. M. Montemore, C. F. Nwaokorie and G. O. Kayode, *Catal. Sci. Technol.*, 2020, **10**, 4467–4476.
- 43 D. Jasrasaria and E. O. Pyzer-Knapp, in *Advances in Intelligent Systems and Computing*, Springer, Cham, 2019, vol. 858, pp. 1–15.
- 44 Ł. Mentel, *Zenodo*, 2021.
- 45 M. M. Bhasin and D. W. Slocum, Eds., *Methane and Alkane Conversion Chemistry*, Springer US, Boston, MA, 1995, vol. 39.
- 46 Y. Meng, C. Ding, Y. Xue, X. Gao, K. Zhang, J. Wang and Z. Li, *New J. Chem.*, 2020, **44**, 3922–3929.
- 47 V. Fung, G. Hu and B. Sumpter, *J. Mater. Chem. A*, 2020, **8**, 6057–6066.
- 48 K. Herrera Delgado, L. Maier, S. Tischer, A. Zellner, H. Stotz and O. Deutschmann, *Catalysts*, 2015, **5**, 871–904.
- 49 G. Jones, J. G. Jakobsen, S. S. Shim, J. Kleis, M. P. Andersson, J. Rossmeisl, F. Abild-Pedersen, T. Bligaard, S. Helveg, B. Hinnemann, J. R. Rostrup-Nielsen, I. Chorkendorff, J. Sehested and J. K. Nørskov, *J. Catal.*, 2008, **259**, 147–160.
- 50 R. T. Hannagan, G. Giannakakis, R. Réocreux, J. Schumann, J. Finzel, Y. Wang, A. Michaelides, P. Deshlahra, P. Christopher, M. Flytzani-Stephanopoulos, M. Stamatakis and E. C. H. Sykes, *Science (80-.)*, 2021, **372**, 1444–1447.
- 51 A. A. Peterson and J. K. Nørskov, *J. Phys. Chem. Lett.*, 2012, **3**, 251–258.
- 52 Y. Pei, H. Zhong and F. Jin, *Energy Sci. Eng.*, 2021, **9**, 1012–1032.
- 53 X. Zhi, Y. Jiao, Y. Zheng, K. Davey and S. Z. Qiao, *J. Mater. Chem. A*, 2021, **9**, 6345–6351.
- 54 D. Behrendt, S. Banerjee, C. Clark and A. M. Rappe, *J. Am. Chem. Soc.*, 2023, **145**, 4730–4735.
- 55 M. T. Darby, R. Réocreux, E. C. H. Sykes, A. Michaelides and M. Stamatakis, *ACS Catal.*, 2018, **8**, 5038–5050.
- 56 J. K. Nørskov, T. Bligaard, A. Logadottir, J. R. Kitchin, J. G. Chen, S. Pandalov and U. Stimming, *J. Electrochem. Soc.*, 2005, **152**, J23.
- 57 H. Guo, H. Li, K. Jarvis, H. Wan, P. Kunal, S. G. Dunning, Y. Liu, G. Henkelman and S. M. Humphrey, *ACS Catal.*, 2018, **8**, 11386–11397.

- 58 S. Tian, B. Wang, W. Gong, Z. He, Q. Xu, W. Chen, Q. Zhang, Y. Zhu, J. Yang, Q. Fu, C. Chen, Y. Bu, L. Gu, X. Sun, H. Zhao, D. Wang and Y. Li, *Nat. Commun.*, 2021, **12**, 3181.