



OPEN

Agents for sequential learning using multiple-fidelity data

Aini Palizhati^{1,2}, Steven B. Torrisi¹, Muratahan Aykol¹, Santosh K. Suram¹, Jens S. Hummelshøj¹ & Joseph H. Montoya¹✉

Sequential learning for materials discovery is a paradigm where a computational agent solicits new data to simultaneously update a model in service of exploration (finding the largest number of materials that meet some criteria) or exploitation (finding materials with an ideal figure of merit). In real-world discovery campaigns, new data acquisition may be costly and an optimal strategy may involve using and acquiring data with different levels of fidelity, such as first-principles calculation to supplement an experiment. In this work, we introduce agents which can operate on multiple data fidelities, and benchmark their performance on an emulated discovery campaign to find materials with desired band gap values. The fidelities of data come from the results of DFT calculations as low fidelity and experimental results as high fidelity. We demonstrate performance gains of agents which incorporate multi-fidelity data in two contexts: either using a large body of low fidelity data as a prior knowledge base or acquiring low fidelity data in-tandem with experimental data. This advance provides a tool that enables materials scientists to test various acquisition and model hyperparameters to maximize the discovery rate of their own multi-fidelity sequential learning campaigns for materials discovery. This may also serve as a reference point for those who are interested in practical strategies that can be used when multiple data sources are available for active or sequential learning campaigns.

A central concern of the materials discovery and optimization process is a simple, practical question: given limited researcher time and resources, what is the next experiment that should be performed? The urgent need for new energy technologies to mitigate fossil fuel use makes this question especially relevant. Widespread adoption of novel fuel cell catalysts, batteries, thermoelectrics, and other energy technologies requires optimization on many different fronts: materials discovery campaigns may target compounds with improved cost, safety, stability, efficacy, or some combination of these and other goals. The use of artificial intelligence tools to accelerate the discovery and optimization process, hand-in-hand with developments in high-throughput experimentation and analysis, may help us to meet timely goals for decarbonization of the global energy economy.

This work is a step towards bridging three relatively recent advances in the materials science research community, which are still realizing their individual and combined potential: (1) the advent of large-scale and freely available databases of computational simulations^{1–4}, particularly from density functional theory (DFT)^{5,6}, (2) the mainstream accessibility of machine learning tools⁷, and (3) development of high-throughput experimentation hardware and software^{8–11}. DFT has shown its applicability in complementing and even guiding the experimental discovery of materials^{12–14}. Machine learning exploits the widespread availability of DFT results to allow accurate and interpretable surrogate models to estimate a desired materials property before either experiment or simulation^{12,15–23}. By increasing the efficiency of theoretical property prediction, the combination of large-scale DFT and machine learning makes it easier for researchers to obtain theoretical predictions for a wider variety of materials, which can then guide the high-throughput experimental process. The paradigm of sequential (or active) learning (henceforth SL), in which a model solicits new training data and updates its performance in response to this data, is useful both to computational^{24–28} and experimental high throughput studies for both optimization and analysis²⁹. Some examples of sequential learning include: systems that learn how to perform only the most valuable or relevant DFT simulations using previous iterations^{25,30,31}, improve force fields more rapidly for molecular dynamics simulations^{24,32}, and synthesize carbon nanotubes at new conditions that promote higher yields and higher qualities of product³³. The sequential learning paradigm thus can provide a conceptual link between materials optimization and discovery workflows across computational and experimental methodologies.

This work also considers a complication in the reality of the scientific discovery process, where there are many sources of data with different costs to obtain them. We use “multi-fidelity” to describe these diverse data

¹Energy and Materials Division, Toyota Research Institute, Los Altos, USA. ²Department of Chemical Engineering, Carnegie Mellon University, Pittsburgh, USA. ✉email: joseph.montoya@tri.global

sources, assuming a tradeoff between data accuracy and cost to obtain. High-fidelity data is considered expensive to obtain but more accurate, whereas low fidelity data is considered cheaper to obtain but less accurate. A promising approach to utilize data with various fidelity is through multi-fidelity models. These models provide a clear conceptual advantage: they can combine large quantities of cheaply acquired, less accurate data with data acquired via more expensive but accurate methods. Multi-fidelity models mitigate resource limitations of building models from high fidelity data and can be significantly more accurate in their predictions than similar models trained on single-fidelity datasets^{34–38}. For example, multi-fidelity models can combine theoretical calculations (such as DFT data) with experimental observations, or compare DFT calculations with a computationally efficient functional like PBE with a more costly but accurate functional like HSE^{39,40} or SCAN^{41,42}. Currently, multi-fidelity modeling methods are typically used and investigated independently from sequential learning in the scientific literature, i.e. the relationship between the acquisition strategy and the modeling strategy is chosen on an ad-hoc basis. This work explores and attempts to show how scientists can use sequential learning and multi-fidelity methods together.

In this work, we introduce a multi-fidelity sequential learning framework for materials discovery based on agents which minimize the number of high fidelity acquisitions while simultaneously optimizing for a figure of merit. These agents are open-source and new additions to the existing software framework CAMD³¹. We demonstrate the agents' capability using electronic band gap data of inorganic compounds at two fidelities: experimental measurements as high fidelity and GGA-level DFT simulations as low fidelity. We choose the band gap as it is a fundamental electronic property relevant for a wide range of technological applications^{14,43,44}. Hence, materials with targeted bandgaps have been the subject of several combinatorial high-throughput efforts including searches for photocatalytic materials with the capability of absorbing sunlight^{45–47}. Thus, there is a large body of experimental and theoretical bandgap data covering a wide range of chemical systems^{35,36,48}. Using our framework and this previously-existing data, we performed several iterative sequential learning campaigns to 'simulate' the discovery of inorganic materials with a target electronic band gap window. The scope of this study is thus to benchmark the process of integrating multiple fidelities in a discovery campaign. This study serves as "wargames" for multi-fidelity sequential learning campaigns so that in applying these tools to a real-world campaign (with its associated costs on researcher time and resources), the agents can be well-chosen for the task with a greater degree of confidence and user knowledge.

We benchmark discovery campaigns in two settings: one where all the low-fidelity or "cheap" data is available from the start, and one where low-fidelity data is acquired in parallel to high-fidelity data. In other words, in the first setting, DFT data is considered prior knowledge, and in the second setting, DFT and experimental data are acquired together in sequential iterations. These sequential learning procedures benchmark how different ML models and acquisition strategies influence the overall rate of discovery of materials per experiment, which was scored using several previously used figures of merit⁴⁹. Our results show that consideration of lower fidelity DFT data in conjunction with experimental data increases the rate of discovery of materials suitable for solar photoabsorption, which suggests that this may be a useful strategy for high-throughput campaigns involving variable data fidelity to use to accelerate their work. We further demonstrate that the type of machine learning model used in the sequential learning procedure controls the extent of an increase in the number of discoveries in the multi-fidelity setting compared to the single-fidelity baseline.

Methods

Below, we will detail how we compiled and encoded the band gap database in "[Dataset collection and representation](#)", how we designed a sequential learning procedure that can use and request data in "[Multi-fidelity sequential learning procedure](#)", how we developed agents which can account for this additional complexity in "[Agent design for materials discovery](#)", how we evaluated the agents in "[Performance metrics for sequential learning](#)", and what our objective for our benchmarking campaigns is in "[Sequential learning objective](#)".

Dataset collection and representation. The band gap dataset was collected from two sources: (1) experimentally reported band gaps of inorganic semiconductors aggregated by Zhuo et al.³⁴ and disseminated via the Matminer⁵⁰ package, and (2) GGA-level DFT-computed band gaps generated and disseminated via the Materials Project database^{1,51}. We first pulled the experimentally reported compositions and their corresponding band gaps. For each experimental composition, we attempted to obtain the band gap corresponding to the most phase-stable (i.e. lowest computed energy per atom) crystal structure from the Materials Project. Here, the DFT-computed band gaps using the Perdew–Burke–Ernzerhof (PBE) functional⁵² were considered low fidelity data, as GGA has well-known systematic errors that underestimate experimentally measured band gaps by ~ 0.9 eV⁵³. Overall, 3960 unique compositions had both experimental and theory data. Out of the 3960 compositions, 375 contained multiple experimental band gap measurements. For each composition with multiple experimental measurements, the respective minimum band gap value was used. Figure 1 describes the dataset by outlining the element occurrence, which shows abundant oxides, sulfides, and selenides, as well as copper and lithium-containing compositions.

We processed the collected data by using a fixed-length vector to encode both the compositions of each material and the level of fidelity. More specifically, stoichiometric compositions were featurized with the matminer ElementProperty featurizer⁵⁰. This featurizer offers flexibility to compare experimental and theoretical data when experimental structure information is not available. The levels of fidelity were represented with one-hot encoding, where a binary variable was added for each fidelity level (i.e. for experimental data, a "1" was placed under the "experiment" feature and a "0" under the "theory" feature and vice versa). To improve numerics, the final overall features were scaled such that their distribution had a mean of 0 and a standard deviation of 1.

H																	He
54																	0
Li	Be											B	C	N	O	F	Ne
320	29											203	103	115	970	104	0
Na	Mg											Al	Si	P	S	Cl	Ar
130	155											191	222	251	476	92	0
K	Ca	Sc	Ti	V	Cr	Mn	Fe	Co	Ni	Cu	Zn	Ga	Ge	As	Se	Br	Kr
167	126	95	153	109	87	180	168	171	204	330	176	189	253	175	425	67	0
Rb	Sr	Y	Zr	Nb	Mo	Tc	Ru	Rh	Pd	Ag	Cd	In	Sn	Sb	Te	I	Xe
114	94	96	96	90	78	23	95	119	142	177	152	233	188	256	241	101	1
Cs	Ba	La	Hf	Ta	W	Re	Os	Ir	Pt	Au	Hg	Tl	Pb	Bi	Po	At	Rn
125	212	122	68	69	67	52	50	95	86	88	126	121	105	159	0	0	0
Fr	Ra	Ac	Rf	Db	Sg	Bh	Hs	Mt	Ds	Rg	Cn	Nh	Fl	Mc	Lv	Ts	Og
0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
Ce	Pr	Nd	Pm	Sm	Eu	Gd	Tb	Dy	Ho	Er	Tm	Yb	Lu				
109	80	96	22	113	92	43	72	81	71	65	68	105	66				
Th	Pa	U	Np	Pu	Am	Cm	Bk	Cf	Es	Fm	Md	No	Lr				
47	0	76	0	0	0	0	0	0	0	0	0	0	0				

Figure 1. The periodic table outlining frequencies of the elements occurring in our band gap dataset. The color represents the frequency of occurrence. This frequency is also labeled under each element symbol.

Multi-fidelity sequential learning procedure. The multi-fidelity sequential learning framework is built on the recently introduced system for Computational Autonomy for Materials Discovery (CAMD)³¹. CAMD is a framework that abstracts decision-making in sequential learning studies into “agents”. Agents perform tasks like training and applying machine learning models or choosing which experiment should be done next based on user-specified criteria. CAMD is open-source and users can add new agents according to their needs (such an agent is one of the contributions of this manuscript). The CAMD framework enables convenient design and testing of acquisition strategies from candidate data points in SL-based optimization.

Figure 2 outlines the CAMD framework and highlights the newly constructed multi-fidelity acquisition feature. In a given series of iterations, termed a *campaign*, the (multi-fidelity) seed data and candidate data (search space) go into an *agent*. A preprocessing step in the agent featurizes each data point in the *seed data* and *candidate data* using the point’s composition and fidelity as described in the previous section. The featurized seed data is used to train a machine learning model, which makes predictions on the candidate data for the target property. Using the predictions, the agent then selects candidates at different fidelities. In the CAMD framework, candidates selected by the agent are sent to an *experiment API*, which collects the experimental data corresponding to the candidate and augments the dataset, allowing candidate data to be moved into the seed data for new active learning iterations. For the sake of active learning simulation, the CAMD experiment is an “after-the-fact” (ATF) API that emulates DFT simulation and experimental measurement, respectively, and returns the results from the known dataset which the agent and CAMD campaign are not aware of prior to the acquisition. This after-the-fact protocol reflects the scope of our study: to benchmark multi-fidelity agent performance in how efficiently they explore a known dataset, demonstrating the gains of multi-fidelity agents with various exploration strategies. The ATF experiment API can be exchanged for one that collects data from and monitors a real experiment, performing new experiments or DFT simulations with the agent that has been designed using ATF simulations³¹. Another CAMD object, the *analyzer*, monitors the campaign results and provides an analysis of the experiments in the context of the previously collected data (i.e. the seed data) and the progress of the campaign. In our case, the analyzer monitors and reports the cumulative number of materials suitable for solar photoabsorption. Upon the completion of the agent selection, experimental acquisition, and analysis phases of a campaign iteration, newly obtained experimental results are appended to the seed data and removed from the candidate data, and campaign begins in a new iteration with agent selection.

Agent design for materials discovery. Designing the agent for a multi-fidelity sequential learning procedure required two steps: (1) selecting appropriate machine learning models and (2) generalizing a CAMD-compatible³¹ data acquisition decision-making process to allow for multiple levels of data fidelity. For model selection, we implemented and compared several well-known regression methods, including support vector regression (SVR), k-nearest neighbors (KNN), random forest regression (RFR), and Gaussian process regression (GPR). For each model, we optimized hyperparameters and did comparative performance analysis (detailed results can be found in Supplementary document S1). Based on the results, support vector regression, random forest regression, and Gaussian process regression had qualitatively similar performances and were used for

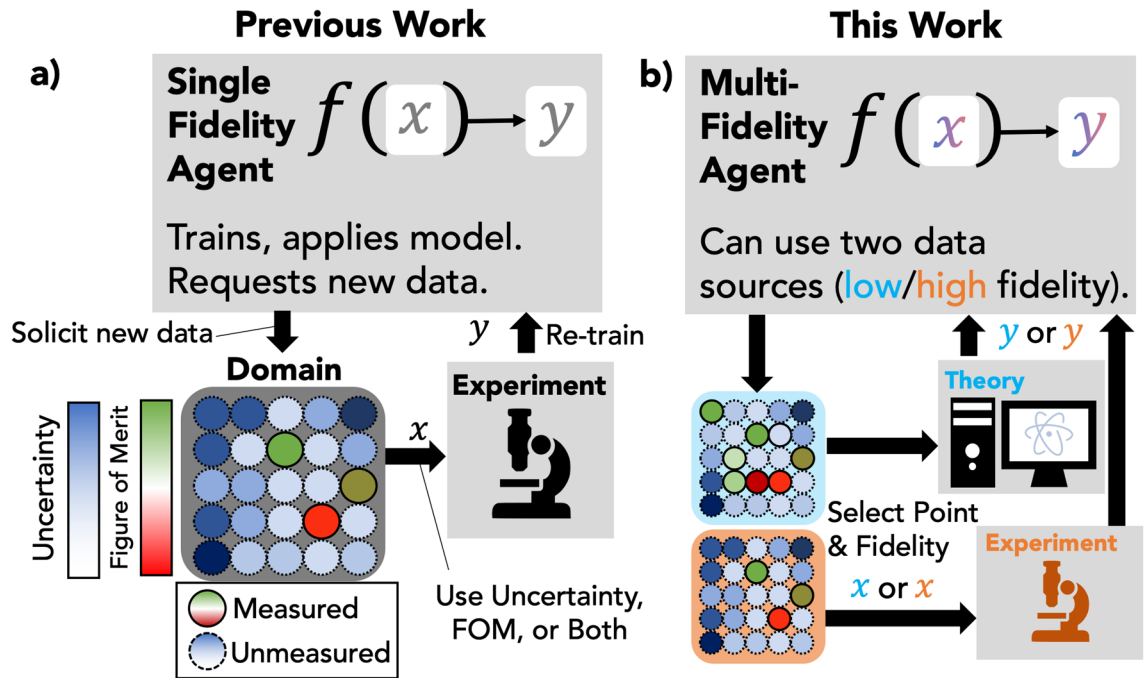


Figure 2. A schematic outlining the distinguishing features of our new multi-fidelity sequential learning framework. (a) Previous work established in Ref³¹. The major components of the framework include an agent, which maintains a model which is iteratively updated in response to new data selected and solicited by the agent. The agent passes the instruction to an “experiment” (a generic term for some kind of data generation, and could alternately be a theoretical calculation). (b) Our work is distinguished by making available two different sources of data with differing degrees of fidelity (here, we use experimental and theoretical data). The agent makes predictions using both sources of data and makes the decision to select new data from one pool or the other.

framework construction and demonstration. Our implementation is sufficiently general to allow users to choose any scikit-learn-compatible ML model and their choice of hyperparameters.

A primary design concern in developing a multi-fidelity agent is mathematically framing the problem of when to draw from low-cost, low-fidelity data vs. high-cost, high-fidelity data. To this end, we designed two agents, an epsilon-greedy multi-fidelity agent (henceforth ϵ -greedy-MF) and a Gaussian process lower confidence bound⁵⁴ derived multi-fidelity agent (GPR_{LCB}-MF). The latter exploits the fact that Gaussian Process regression allows for a principled uncertainty estimate “out-of-the-box”, whereas the former works for regression algorithms lacking this feature.

The salient features of the ϵ -greedy agent are that it takes as input a budget of high-fidelity datapoints n which controls the balance between low-fidelity and high-fidelity data. The agent will only call for high-fidelity measurements in domains that have been previously covered by low-fidelity data (see details in Algorithm S1). The ϵ -greedy-MF agent works using any supervised machine learning regressor from scikit-learn⁷ as input. Meanwhile, the GPR_{LCB}-MF agent operates under a total acquisition budget and calls for low- or high-fidelity data in a more sophisticated way. It acquires candidates factoring in Gaussian process regression predicted uncertainties in the LCB setting and *hallucination* of information gain from low fidelity acquisitions analogous to work of Desautels et al. in batch mode LCB⁵⁵ (see full details in Algorithm S2). *Hallucination* works as such: for a high fidelity candidate, the GPR_{LCB}-MF agent adds the lower fidelity predicted posterior mean into the seed data. As a consequence, the higher fidelity candidate prediction gets updated. Essentially, hallucination refers to the ability of the agent to predict ahead of time how low-fidelity data will impact the uncertainty estimate of the model. Hallucination allows the agent to use low-fidelity candidates to explore potentially promising parts of the domain, while using high-fidelity candidates to exploit promising regions of parameter space, offloading exploratory (higher risk) acquisitions first to lower-fidelity computations. In our formulation, three hyperparameters that must be empirically optimized govern the tradeoff between data fidelities: α , β , and γ . α is the uncertainty multiplier in GPR_{LCB} as shown below:

$$\hat{y}_{i,LCB} = \hat{y}_i - \alpha * \sigma_i \quad (1)$$

where \hat{y}_i is the posterior mean and σ_i is the uncertainty given a candidate i . α here sets the weight of uncertainty in the LCB setting. Next, β is a threshold for uncertainty. For a given observation, if its σ_i is less than β , the observation is considered to have low uncertainty. A small β makes the agent “risk-averse” around high-fidelity measurements in unexplored regions of space. In the small β regime, unless the uncertainty on a given prediction is very low, it will acquire lower fidelity data first. Inversely, if β is large, the agent is tolerant to high uncertainty for experiments and will more readily add experimental data. In practical applications, β could be set with respect

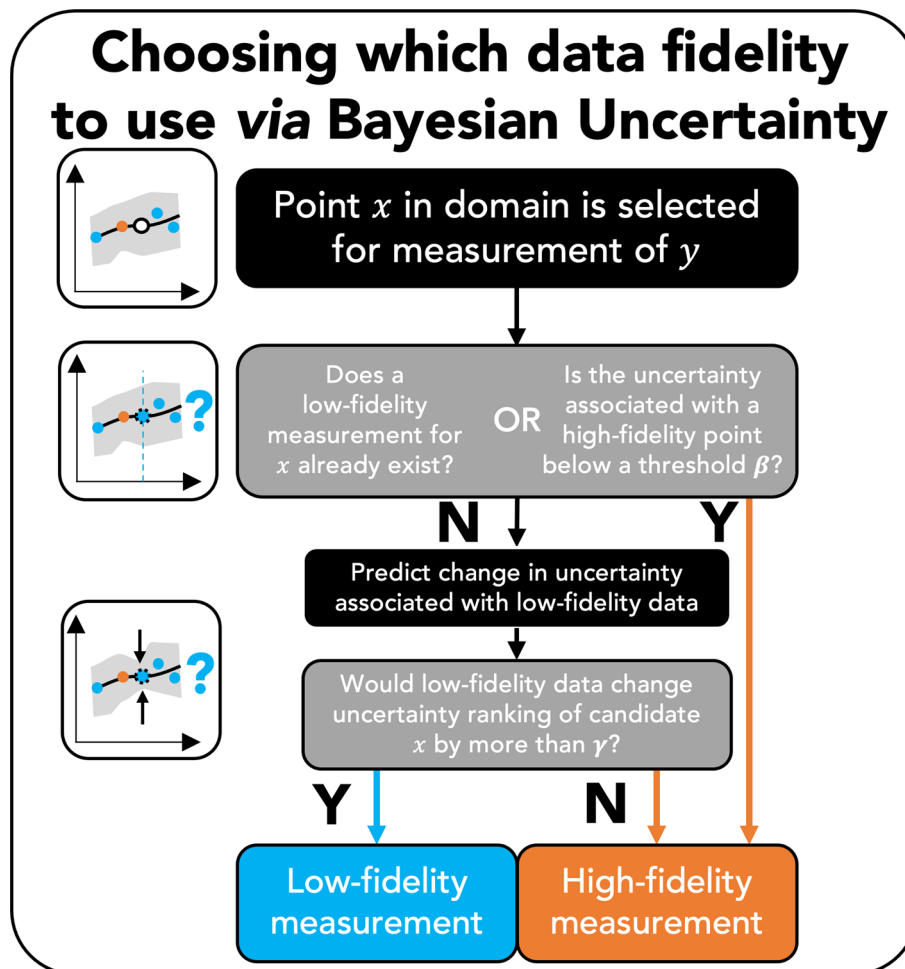


Figure 3. Overview of how uncertainty is used to balance the use of low and high fidelity data in the implementation of Algorithm 2 of Section S2. When a data point x is called for as an experimental candidate, the above flowchart describes the decision-making process for which data source to use. After a data point x is selected for measurement, two conditions are checked: (1) if the corresponding (e.g. with the same formula) low-fidelity measurement has already been made or (2) if the uncertainty associated with the high fidelity measurement is low enough (below a threshold β). If either is true, the high-fidelity measurement is taken. If neither are true, then the agent must consider the trade-off between low-fidelity and high-fidelity data. The agent thus considers how a low fidelity data point would affect the current ordering of the predicted figure of merit associated with all candidates. If it would alter the ranking by more than γ , the low-fidelity measurement is taken. If not, the high-fidelity measurement is taken. Note that β and γ are user-defined hyperparameters explained in detail in Section S2.

to the cost of acquiring high-fidelity data. Lastly, γ is a threshold for the influence of hallucination (denote Δr) as shown below:

$$\Delta r = r_i^* - r_i \quad (2)$$

Here, r_i is the ranking of an observation $\hat{y}_{i,LCB}$ in the candidate space based on its distance to the target value. r_i^* is the new ranking of the observation after hallucination. The agents acquire a high fidelity candidate if $\Delta r \leq \gamma$. If γ is 0, then the prospect of the lower fidelity data has to increase the chances of the experiment being successful. Otherwise, lower fidelity data will be acquired first. Because in this case, r_i^* has to be a smaller value than r_i (i.e. a better ranking). If γ is very large, then the agent does not care about how much low fidelity data affects the potential experiment. Because in this case, r_i^* can be any value, including a value that is higher than r_i (i.e. a worse ranking). The overall influence of these hyperparameters is summed up in a broad overview way in Fig. 3. We simulated various scenarios of these three hyperparameters to optimize the agents. The details and results are in Supplementary document section S3. The Gaussian processes were implemented using the GPy⁵⁶ package. Details of agents are also made explicit in the code available via the open-source CAMD repository at <https://github.com/TRI-AMDD/CAMD>.

Performance metrics for sequential learning. To quantitatively compare the efficacy (the ability to call for data points that optimize a figure of merit) and efficiency (the number of experiments required to do so) between active learning campaigns, we used previously established active learning metrics (ALM) described by Rohr et al.⁴⁹. These metrics are ALM, acceleration factor (AF), and enhancement factor (EF), defined as follows and explicated further below:

$$ALM(x, N_{exp}) = \frac{\text{Target Materials Discovered after } N_{exp}}{\text{Total Target Materials}} \quad (3)$$

$$AF(x, y, ALM) = N_{exp}(x|ALM) - N_{exp}(y|ALM) \quad (4)$$

$$EF(x, y, N_{exp}) = ALM(x, N_{exp}) - ALM(y, N_{exp}) \quad (5)$$

where x and y are agents, N_{exp} is the number of experiments performed (i.e. in our case, the high fidelity data acquired). The function N_{exp} conditioned on ALM (i.e. $N_{exp}(x|ALM)$ and $N_{exp}(y|ALM)$) refers to the number of experiments in the sequential learning campaigns that attained an ALM. Because our emulated discovery campaigns are trying to find materials with a visible-spectrum band-gap, our discovery process can be scored in a binary way: any new data point's band gap is either inside or outside of the target range. Thus, we can compute the fraction of ideal materials which were correctly identified at an iteration given a sequential learning run and so ALM lies within $[0, 1]$. This metric is defined for a single sequential learning campaign and is most useful for after-the-fact workflows, as the denominator requires some knowledge of the total number of materials which lie within the target range (For a 'real-world' case where the materials are not known ahead of time, the final number of target materials discovered by the campaign can be used in scoring sequential learning, as ALM is a 'time-dependent' property that can change at each iteration step. Also, note that this study focuses on materials that are scored in a binary way as having the band gap property within a target range. For cases where a quantity is optimized around some target, this could be defined using the distance of the best-known material thus far to the current best-known global maximum/minimum target property). Next, acceleration factor and enhancement factor are metrics that compare *two* sequential learning runs to one another using the ALM. The acceleration factor is the reduction of required budget (e.g. in time, iterations, or some other consumed resource) between an agent and a benchmark case (e.g. random selection, an alternate model, single-fidelity, or manual human selection) to reach a particular fraction of ideal candidates ($AF = N_{budget,benchmark} - N_{budget,agent}$). In other words, given ALM vs. N_{exp} , the acceleration factor is the "horizontal-line" distance between two models at an ALM at different "times". A positive value of acceleration factor between a multi-fidelity campaign and a single-fidelity campaign means the former outperformed the latter because it reduced the required budget needed to achieve a certain amount of discovery. Similarly, the enhancement factor is the "vertical-line" distance between two campaigns' ALM score at a given "time", which shows the performance enhancement at the same consumed experiment budget. More specifically, at the same number of iterations, amount of elapsed time, or some other metric of expended resources, enhancement factor quantifies the improvement of materials discovery by a given sequential learning method versus a benchmark method ($EF = \frac{N_{discovery,agent}}{N_{discovery,benchmark}}$). In the case of comparing a multi-fidelity campaign to a single-fidelity campaign, when the enhancement factor is greater than one, it indicates that the multi-fidelity campaign outperforms its corresponding single-fidelity campaign at a given budget.

Sequential learning objective. For multi-fidelity sequential learning campaign simulations and subsequent performance evaluations of the agents, we attempted to model a discovery campaign for photoabsorbers by targeting materials with experimentally measured band gap $\subseteq [1.6, 2.0]$ eV⁵⁷, i.e. those with reasonable solar photoabsorption, were considered ideal and set as the targets. 207 of our 3960 candidate experimental materials are considered ideal based on the target band gap window defined above. In other words, only about one in twenty or 5% of the candidate materials lie within the target window of the discovery campaign.

Results

Figure 4 highlights the campaigns that we performed to benchmark the sequential learning models. In "Boundary cases: all or no DFT data available" (corresponds to campaign A), we demonstrate the performance gains which come from an agent with full a priori knowledge of DFT calculations soliciting experimental data versus an agent which never uses DFT data exploring the same space of experiments. Next, in "In-tandem acquisitions: both DFT and experiment data are acquired" (corresponds to campaign B), we compare the performance of agents seeded with first 500 experimentally discovered compositions in a multi-fidelity versus single-fidelity context, where either both DFT and experimental data are solicited in-tandem (with some DFT data supplied a priori) or exclusively experimental data seeded and solicited. In both cases, we find that the performance of multi-fidelity agents are improved by the inclusion of low-fidelity DFT data.

Boundary cases: all or no DFT data available. We first tested the acquisition performance of multi-fidelity agents in the limiting case where the full suite of DFT calculations was considered as a priori knowledge. The objective here was to determine how an automated experimental sequential learning procedure would be enhanced by a priori knowledge of a large theoretical dataset. This type of acquisition is for a use case where low-fidelity experimental data is much cheaper to acquire than high-fidelity data that full domain coverage is available at the outset of a high-fidelity experimental campaign. Because no new low-fidelity data is solicited,

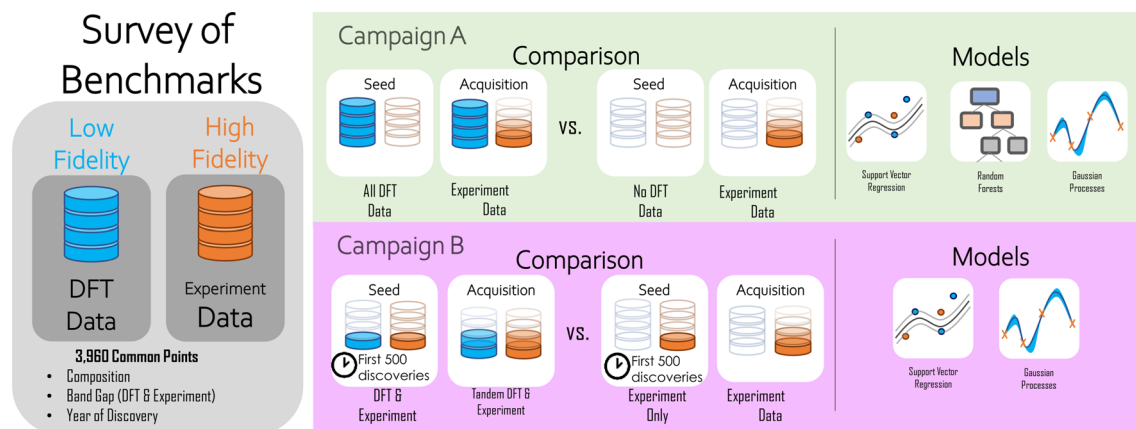


Figure 4. Campaigns performed for single-fidelity acquisition and multi-fidelity acquisition. In campaign A, we compare the results of including all or no DFT data in the seed data set and only acquire experimental data. This campaign effectively compares an “a priori” agents to a single-fidelity agents. In campaign B, the multi-fidelity agents are seeded with first 500 experimentally discovered compositions (based on ICSD⁵⁸ timeline of their first publication³⁹) and their corresponding DFT data. Both DFT and experimental data are acquired here. The single-fidelity agents are exclusively seeded with and acquire experimental data.

gains in campaign performance are entirely due to the transfer of knowledge from the large, low-fidelity dataset in making predictions and subsequent acquisitions under the high-fidelity, expensive setting.

We performed after-the-fact discovery runs with three agents: ϵ -greedy agents that used support vector regression and random forest regression, and a GPR_{LCB} agent. As mentioned previously, ϵ -greedy agents works for regression models lack principled uncertainty estimate, and GPR agent acquire candidates based on both the predicted posterior mean and uncertainty from Gaussian process regression. For each agent, we considered two cases: (1) no low-fidelity seed data at any point in the campaign and (2) all available DFT data as seed data at the outset. Note that both (1) and (2) are thus only acquiring high-fidelity data, and this set of six campaigns benchmarks in the most extreme case if and how much a priori low fidelity knowledge can assist in the discovery campaign. For convenience, we designate $\text{SVR-SF}_{boundary}$, $\text{RFR-SF}_{boundary}$, and $\text{GPR}_{LCB-SF}_{boundary}$, $\text{SVR-MF}_{boundary}$, $\text{RFR-MF}_{boundary}$, and $\text{GPR}_{LCB-MF}_{boundary}$ (SF denotes single-fidelity, MF denotes multi-fidelity). We gave all the agents a budget of 20 experiment requests in each iteration and simulated each campaign for 100 iterations. In addition, several campaigns have additional stochasticity that requires some thought. More specifically, single-fidelity campaigns with no seed data (i.e. $\text{SVR-SF}_{boundary}$, $\text{RFR-SF}_{boundary}$, $\text{GPR}_{LCB-SF}_{boundary}$) create initial seeds data randomly, and random forests also have randomness during the bootstrapping of the samples used in building trees. Even though this stochasticity does not change the candidate acquisition strategy of the agents, it could result in varied campaign performance depending on the inputted random seeds. To account for this, we performed ten trials of campaigns that used those four agents (i.e. all three single-fidelity agents and $\text{RFR-SF}_{boundary}$). This helps us look at the overall campaign performance of those agents more objectively because we have better information about the “average” and “variance” in the performance.

Lastly, we bound the performance of our agents above and below by two limiting cases: (1) a perfect agent, where every acquisition is an ideal candidate and the full target space is explored in exactly 202 steps and (2) a naive agent that chooses the next data point from the candidate space at random.

Figure 5 shows the results of the simulated discovery campaigns. Where the fraction of the target materials found is plotted against the number of experiments (i.e. high fidelity candidate acquired). The shaded region is the standard deviation of materials found for campaigns with multiple trials. For our initial benchmark, we primarily compare the performance between models. In the single-fidelity case with no access to low-fidelity DFT data, looking at the average target materials found (colored dash lines) in each campaign, random forests agent outperformed support vector regression and Gaussian process regression agents until ~ 850 experiment requests, at which point close to 60% of the ideal candidates had been discovered. Support vector regression agent started outperforming the other two from ~ 850 experiment requests. In the multi-fidelity case where all low-fidelity (DFT) data was made available (colored solid lines), all agents performed similarly (with random forests slightly ahead) until ~ 550 experiment requests. Afterward, the support vector regression and Gaussian process regression agent outperformed the rest until the end. More importantly, we observed that multi-fidelity agents outperformed their single-fidelity counterparts, demonstrating that these regression algorithms can transfer the knowledge available from the lower-fidelity dataset in making predictions for the high-fidelity target. All of our sequential learning agents consistently outperformed random acquisitions.

To compare the performance of single and multi-fidelity agents in more detail, we tabulated acceleration factors at 50% and 80% of the total discovery of target candidates in Table 1. To achieve the discovery of 50% of the candidates designed as ideal, multi-fidelity agents reduce the experiments requested by 160, 80, and 180 for support vector regression, random forests, and Gaussian process regression, respectively. At 80% discovery, the acceleration factors are 160, 60, and 220 for support vector regression, random forests, and Gaussian process regression agents, respectively. The enhancement factors shown in Fig. 6 provided a clearer picture of the comparative performance throughout the campaign. We observe that support vector regression multi-fidelity

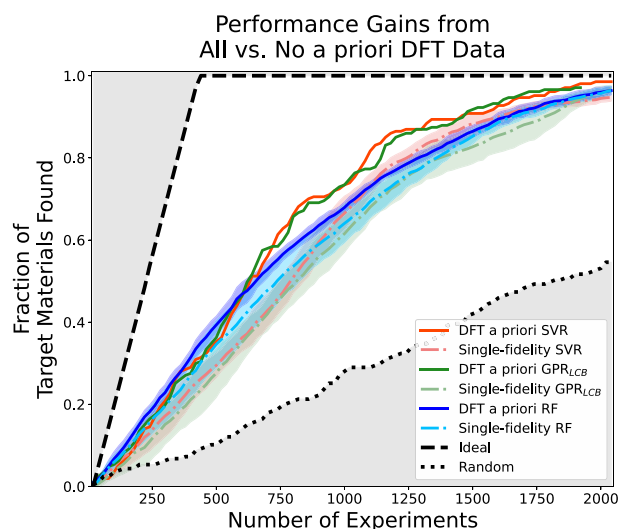


Figure 5. Simulated performance of various campaigns with and without supplying the lower fidelity (DFT) data in the seed. The x-axis corresponds to the high-fidelity (experiments) acquisition budget. The y-axis corresponds to the fraction of ideal materials discovered from the search space. SVR, RF, and GPR_{LCB} correspond to agents using support vector regression, random forests, and Gaussian process regression lower confidence bound, respectively. The shaded colored regions are the standard deviation of materials found for campaigns with multiple trials. The random acquisition and ideal acquisition baselines are also labeled in the figure, representing the lower and upper bounds of agent performance.

Agents	Single-fidelity experiments performed for 50% discovery	Multi-fidelity $AF_{50\%}$	Single-fidelity experiments performed for 90% discovery	Multi-fidelity $AF_{90\%}$
Support vector regression	800	160	1380	160
Random forest regression	740	80	1360	60
GPR_{LCB}	820	180	1360	220

Table 1. Acceleration factor (AF) of multi-fidelity agents in simple acquisitions. The AFs are the reduction in number of experiments performed by multi-fidelity agents to achieve a certain amount of discoveries. For each row, we highlighted the agents used, the experiments performed by single-fidelity agents to achieve 50% and 80% discovery, and the acceleration factor of the multi-fidelity agents at those discoveries.

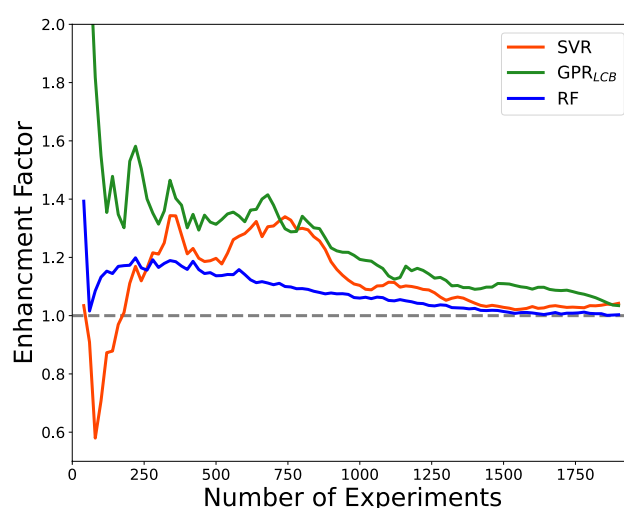


Figure 6. The enhancement factor for quantifying the added value of DFT in the campaigns. The enhancement factors are the ratio of the fraction of discovered materials between the DFT a priori and single-fidelity models at any given budget.

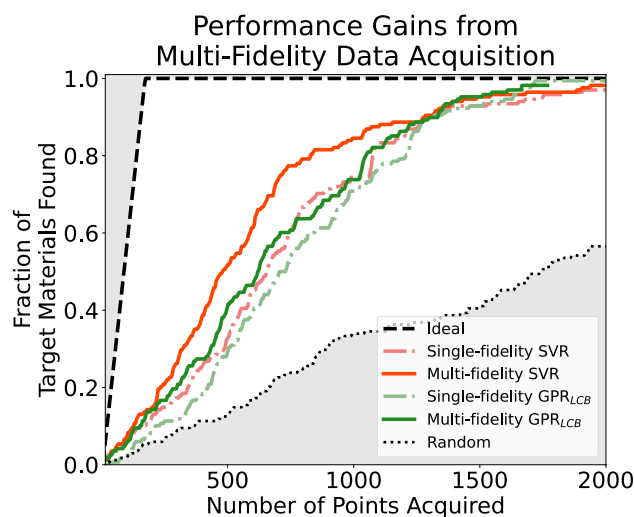


Figure 7. Simulated performance of various campaigns when both DFT and experiment data were acquired during runs. The results of all campaigns are combined in one figure. $N_{\text{experiment}}$ on the x-axis corresponds to the experimental acquisition budget. Discoveries on the y-axis correspond to the fraction of ideal materials discovered from the search space. The random acquisition and ideal acquisition baselines are also labeled in the figures.

agents briefly underperformed their single-fidelity counterparts in the early stages of campaigns (until ~ 100 experiments). After this point, SVR-MF_{boundary} outperformed SVR-SF_{boundary} by a notable margin to achieve enhancement of a factor of ~ 1.2 to 1.4 until ~ 1000 experiments. This factor diminished slowly as candidates were exhausted for the remainder of the campaign. GPR_{LCB}-MF_{boundary} and RFR-MF_{boundary} consistently outperformed their single-fidelity counterpart, with GPR_{LCB}-MF_{boundary} having larger enhancement factors. We also notice a similar diminishing trend of their enhancement factors as candidates were exhausted. In summary, all multi-fidelity agents outperformed their single-fidelity counterparts at all points in the process until most target candidates have been acquired. Between the three agents used, support vector regression and Gaussian process regression agents benefited more from a priori data based on the metrics computed.

In-tandem acquisitions: both DFT and experiment data are acquired. Having investigated two boundary scenarios in the previous section, with all-or-no low fidelity data, we now turn to our next main question: when and how should we decide to acquire low-fidelity data to support and minimize the number of high-fidelity measurements *during* a sequential, closed-loop data acquisition procedure? To answer this, we simulated another set of campaigns benchmarking single-fidelity versus multi-fidelity. First, to mimic a more true-to-life discovery process, we split the compositions into seed data and candidate data based on their year of discovery according to the ICSD⁵⁸ timeline of their first publication⁵⁹ (Fig. 4). In other words, this rationale for selecting the seed data makes the initial data used for the runs and the successive choice of data by the models entirely deterministic. For single-fidelity campaigns, the data of the first 500 experimentally discovered compositions, up to the discovery year of 1965, were included in the seed data, the remaining (3460 compositions) were included in the candidate data. For multi-fidelity campaigns, the data split was identical, with the addition of corresponding DFT data in each set. Next, we set up the campaigns with a ϵ -greedy agent that used support vectors and a Gaussian processes regression agent (since these two agents had better gains in “Boundary cases: all or no DFT data available”). Therefore, a total of four campaigns were set up: SVR-SF_{tandem}, GPR_{LCB}-SF_{tandem}, SVR-MF_{tandem}, and GPR_{LCB}-MF_{tandem} (SF denotes single-fidelity, MF denotes multi-fidelity). As before, we also included the two limiting cases of (1) random acquisition and (2) ‘perfect’ acquisition. For the acquisition budget, both SVR-SF_{tandem}, GPR_{LCB}-SF_{tandem}, along with the two limiting cases, had a budget of 5 experiment requests. SVR-MF_{tandem} had a fix-ratio budget of 5 experiments and 5 DFT. GPR_{LCB}-MF_{tandem} had a budget of 5 acquisitions, each acquisition can be either experiments or DFT, depending on the uncertainties and *hallucination* of information gained from DFT. Based on optimization results in Supplementary document section S3, $\alpha=0.08$, $\beta=5$, and $\gamma=10$ were used for GPR_{LCB}-MF_{tandem} to compare against the other sequential learning campaigns. All campaigns were run until 2000 experiments have been acquired, unless it is stopped due to no discovery after 30 iterations (a setting in the campaign hyperparameter).

Figure 7 shows the qualitative results of the simulated campaigns using in-tandem acquisition. Here, SVR-MF_{tandem} agent outperformed its single-fidelity counterpart early in the campaigns (when $N_{\text{experiments}}$ reached ~ 100). It then stayed ahead until ~ 1200 experiments were acquired, at which point 90% of the ideal materials had been discovered. GPR_{LCB}-MF_{tandem} agent also outperformed its single-fidelity counterpart until 90% of the ideal materials have been discovered (at ~ 1200 experiments). Compared among all four agents, SVR-MF_{tandem} agent’s performance was the best. Furthermore, GPR_{LCB}-MF_{tandem} agent’s performance was similar to that of SVR-SF_{tandem}’s.

Agents	Single-fidelity experiments performed for 50% discovery	Multi-fidelity AF _{50%}	Single-fidelity experiments performed for 80% discovery	Multi-fidelity AF _{80%}
Support vector regression	655	175	1080	250
GPR _{LCB}	705	85	1210	159

Table 2. Acceleration factor of multi-fidelity agents in in-tandem acquisitions. For each row, we highlighted the agents used, the experiments performed by single-fidelity agents to achieve 50% and 80% discovery, and the acceleration factor (AF) of the multi-fidelity agents. The AF's are the reduction in number of experiments performed.

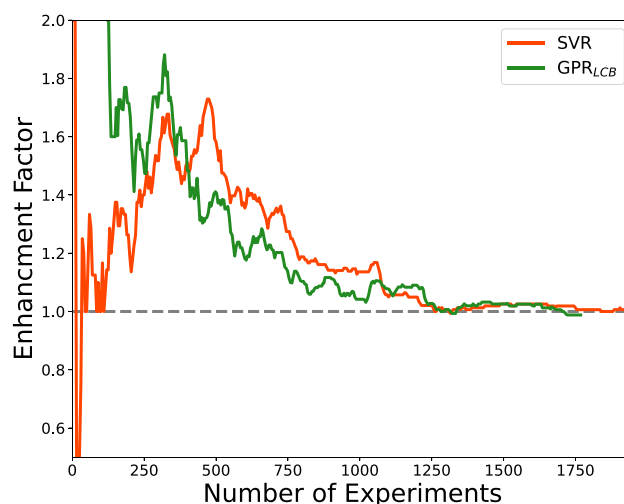


Figure 8. The enhancement factors of in-tandem acquisitions as the campaign progresses.

The acceleration factor (Table 2) of multi-fidelity acquisitions at 50% discovery were 175 and 85 for in-tandem support vector machines and Gaussian processes respectively. At 80% discovery, they were 250 and 159, respectively. The enhancement factors (Fig. 8) of in-tandem multi-fidelity support vector regression is very noisy at first (until $N_{\text{experiments}}$ reached ~ 150), which agrees with Fig. 7. Then they stayed above 1 until $N_{\text{experiments}}$ reached ~ 1250 . The enhancement of GPR_{LCB}-MF_{tandem} cannot be calculated at first because its single-fidelity counterpart did not have any discovery. After $N_{\text{experiments}}$ reached ~ 100 and its single-fidelity counterpart made some discoveries, the enhancement factors were high but decreased as the acquisition continued and converged to 1 at ~ 1250 experiments.

Conclusion

In this work, we develop, implement, and benchmark sequential learning agents that allow for the differentiation of data points of different fidelities. Using our implementation in the CAMD sequential learning framework, we simulated a materials discovery process on previously existing experimental and theoretical electronic band-gap data to inform the selection of these models and suggest hyperparameters that could be used to accompany a ‘real-life’ data acquisition campaign. We found that when all low-fidelity data were provided as a priori knowledge, all multi-fidelity agents outperformed their single-fidelity counterparts and sustained a materials discovery acceleration of 20–60% early on in the campaigns. As the number of experiments acquired in the seed data increased, we saw a decline in additional gain for those multi-fidelity agents. When acquiring low and high-fidelity data in-tandem with support vector regression and Gaussian process regression multi-fidelity agents, both of them still outperformed their single-fidelity counterparts, suggesting strategic acquisitions of lower fidelity data provides a transfer of knowledge and augment higher fidelity target material discovery. We note that, Gaussian process regression multi-fidelity agent here barely outperformed support vector regression single fidelity agent with the settings we provided, which suggests further investigations of the agent.

In summary, we observed a clear trend of multi-fidelity sequential learning agents outperforming those which may only sample at a single-fidelity. The results demonstrate that for studies where low-fidelity data is extremely cheap relative to high-fidelity data, the inclusion of separately labeled data either “up-front” or acquired in-tandem with high-fidelity experiments can increase the rate at which valuable experiments are performed. However, the relative performance of multi-fidelity acquisition is sensitive to the dataset size, ML model selection, and acquisition strategy. Furthermore, the multi-fidelity agents can be extended to have data inputs beyond two fidelities. As mentioned in “Dataset collection and representation”, since the level of fidelity is represented with one-hot encoding, additional fidelity can be passed as an additional column in the feature. Subsequently, acquisition strategies are easily adaptable for multiple levels of fidelity for both of our proposed

algorithms by replicating the logic in a nested fashion. For example, for three levels of fidelity, one may acquire the lowest level of fidelity in order to reduce the uncertainty on a median level of fidelity, and acquire a median level of fidelity to reduce the uncertainty of the highest level of fidelity until the experimental budget threshold for new experiments is reached.

Given these dependencies, our framework offers a critical capability that frames automated discovery process itself as an object of study. Our study on multi-fidelity sequential learning campaigns lays a foundation for future research in which both simulations and experiments can be conducted in-tandem with strategies optimized for their relative cost and accuracy.

Code availability

The full details of the code are provided in an open-source repository at <https://github.com/TRI-AMDD/CAMD>.

Received: 19 November 2021; Accepted: 17 February 2022

Published online: 18 March 2022

References

- Jain, A. *et al.* Commentary: The materials project: A materials genome approach to accelerating materials innovation. *APL Mater.* **1**, 011002. <https://doi.org/10.1063/1.4812323> (2013).
- Kirklin, S. *et al.* The open quantum materials database (OQMD): Assessing the accuracy of DFT formation energies. *NPJ Comput. Mater.* **1**, 1–15. <https://doi.org/10.1038/npjcompumats.2015.10> (2015).
- Curtarolo, S. *et al.* AFLOW: An automatic framework for high-throughput materials discovery. *Comput. Mater. Sci.* **58**, 218–226. <https://doi.org/10.1016/j.commatsci.2012.02.005> (2012).
- Ong, S. P. Accelerating materials science with high-throughput computations and machine learning. *Comput. Mater. Sci.* **161**, 143–150. <https://doi.org/10.1016/j.COMMATSCI.2019.01.013> (2019).
- Hohenberg, P. & Kohn, W. Inhomogeneous electron gas. *Phys. Rev.* **136**, B864–B871. <https://doi.org/10.1103/PhysRev.136.B864> (1964).
- Kohn, W. & Sham, L. J. Self-consistent equations including exchange and correlation effects. *Phys. Rev.* **140**, A1133–A1138. <https://doi.org/10.1103/PhysRev.140.A1133> (1965).
- Pedregosa, F. *et al.* Scikit-learn: Machine learning in Python. *J. Mach. Learn. Res.* **12**, 2825–2830 (2011).
- Hattrick-Simpers, J. R., Gregoire, J. M. & Kusne, A. G. Perspective: Composition-structure-property mapping in high-throughput experiments: Turning data into knowledge. *APL Mater.* **4**, 053211. <https://doi.org/10.1063/1.4950995> (2016).
- Stach, E. *et al.* Autonomous experimentation systems for materials development: A community perspective. *Matter* **4**, 2702–2726. <https://doi.org/10.1016/j.MATT.2021.06.036> (2021).
- Roch, L. M. *et al.* ChemOS: An orchestration software to democratize autonomous discovery. *PLoS One* **15**, e0229862. <https://doi.org/10.1371/JOURNAL.PONE.0229862> (2020).
- Al Hasan, N. M. *et al.* Combinatorial exploration and mapping of phase transformation in a ni-ti-co thin film library. *ACS Combin. Sci.* **22**, 641–648. <https://doi.org/10.1021/acscombsci.0c00097> (2020).
- Schmidt, J., Marques, M. R., Botti, S. & Marques, M. A. Recent advances and applications of machine learning in solid-state materials science. *NPJ Comput. Mater.* <https://doi.org/10.1038/s41524-019-0221-0> (2019).
- Cai, J., Chu, X., Xu, K., Li, H. & Wei, J. Machine learning-driven new material discovery. *Nanosc. Adv.* **2**, 3115–3130. <https://doi.org/10.1039/d0na00388c> (2020).
- Jain, A., Shin, Y. & Persson, K. A. Computational predictions of energy materials using density functional theory. *Nat. Rev. Mater.* **1**, 1–13. <https://doi.org/10.1038/natrevmats.2015.4> (2016).
- Tran, K., Palizhati, A., Back, S. & Ulissi, Z. W. Dynamic workflows for routine materials discovery in surface science. *J. Chem. Inf. Model.* **58**, 2392–2400. <https://doi.org/10.1021/ACS.JCIM.8B00386> (2018).
- Gu, G. H., Noh, J., Kim, I. & Jung, Y. Machine learning for renewable energy materials. *J. Mater. Chem. A* **7**, 17096–17117. <https://doi.org/10.1039/c9ta02356a> (2019).
- Dan, Y. *et al.* Generative adversarial networks (GAN) based efficient sampling of chemical composition space for inverse design of inorganic materials. *NPJCM* **6**, 84. <https://doi.org/10.1038/S41524-020-00352-0> (2020). (arXiv:1911.05020).
- Erdem Günay, M. & Yıldırım, R. Recent advances in knowledge discovery for heterogeneous catalysis using machine learning. *Catal. Rev. Sci. Eng.* <https://doi.org/10.1080/01614940.2020.1770402> (2020).
- Jennings, P. C., Lysgaard, S., Hummelshøj, J. S., Vegge, T. & Bligaard, T. Genetic algorithms for computational materials discovery accelerated by machine learning. *NPJ Comput. Mater.* **5**, 46. <https://doi.org/10.1038/s41524-019-0181-4> (2019).
- Ward, L. *et al.* Including crystal structure attributes in machine learning models of formation energies via Voronoi tessellations. *Phys. Rev. B* **96**, 024104. <https://doi.org/10.1103/PhysRevB.96.024104> (2017).
- Xie, T. & Grossman, J. C. Crystal graph convolutional neural networks for an accurate and interpretable prediction of material properties. *Phys. Rev. Lett.* <https://doi.org/10.1103/PhysRevLett.120.145301> (2018).
- Palizhati, A., Zhong, W., Tran, K., Back, S. & Ulissi, Z. W. Toward predicting intermetallics surface properties with high-throughput DFT and convolutional neural networks. *J. Chem. Inf. Model.* <https://doi.org/10.1021/acs.jcim.9b00550> (2019).
- Torrissi, S. B. *et al.* Random forest machine learning models for interpretable x-ray absorption near-edge structure spectrum-property relationships. *NPJ Comput. Mater.* **6**, 109. <https://doi.org/10.1038/s41524-020-00376-6> (2020).
- Vandermause, J. *et al.* On-the-fly active learning of interpretable bayesian force fields for atomistic rare events. *NPJ Comput. Mater.* **6**, 20. <https://doi.org/10.1038/s41524-020-0283-z> (2020).
- Tran, K. & Ulissi, Z. W. Active learning across intermetallics to guide discovery of electrocatalysts for CO₂ reduction and H₂ evolution. *Nat. Catal.* **1**, 696–703. <https://doi.org/10.1038/s41929-018-0142-1> (2018).
- Tian, Y., Lookman, T. & Xue, D. Efficient sampling for decision making in materials discovery. *Chin. Phys. B* **30**, 050705. <https://doi.org/10.1088/1674-1056/ABF12D> (2021).
- Kusne, A. G. *et al.* On-the-fly closed-loop materials discovery via Bayesian active learning. *Nat. Commun.* **11**, 1–11. <https://doi.org/10.1038/s41467-020-19597-w> (2020).
- Bassman, L. *et al.* Active learning for accelerated design of layered materials. *NPJ Comput. Mater.* **4**, 1–9. <https://doi.org/10.1038/s41524-018-0129-0> (2018).
- Noack, M. M. *et al.* Gaussian processes for autonomous data acquisition at large-scale synchrotron and neutron scattering facilities. *Nat. Rev. Phys.* <https://doi.org/10.1038/s42254-021-00345-y> (2021).
- Seko, A. *et al.* Prediction of low-thermal-conductivity compounds with first-principles anharmonic lattice-dynamics calculations and bayesian optimization. *Phys. Rev. Lett.* **115**, 205901. <https://doi.org/10.1103/PhysRevLett.115.205901> (2015).
- Montoya, J. H. *et al.* Autonomous intelligent agents for accelerated materials discovery. *Chem. Sci.* <https://doi.org/10.1039/D0SC01101K> (2020).

32. Coley, C. W., Eyke, N. S. & Jensen, K. F. Autonomous discovery in the chemical sciences Part I: Progress. *Angew. Chem. Int. Ed.* <https://doi.org/10.1002/anie.201909987> (2020).
33. Nikolaev, P. *et al.* Autonomy in materials research: A case study in carbon nanotube growth. *NPJ Comput. Mater.* **2**, 1–6. <https://doi.org/10.1038/npjcompumats.2016.31> (2016).
34. Zhuo, Y., MansouriTehrani, A. & Brgoch, J. Predicting the band gaps of inorganic solids by machine learning. *J. Phys. Chem. Lett.* **9**, 1668–1673. <https://doi.org/10.1021/acs.jpclett.8b00124> (2018).
35. Pilaian, G., Gubernatis, J. E. & Lookman, T. Multi-fidelity machine learning models for accurate bandgap predictions of solids. *Comput. Mater. Sci.* **129**, 156–163. <https://doi.org/10.1016/j.commatsci.2016.12.004> (2017).
36. Chen, C., Zuo, Y., Ye, W., Li, X. & Ong, S. P. Learning properties of ordered and disordered materials from multi-fidelity data. *Nat. Comput. Sci.* **1**, 46–53. <https://doi.org/10.1038/s43588-020-00002-x> (2021).
37. Kandasamy, K., Dasarathy, G., Schneider, J. & Póczos, B. Multi-fidelity Bayesian optimisation with continuous approximations. In *Proceedings of the 34th International Conference on Machine Learning, vol 70 of Proceedings of Machine Learning Research* (eds Precup, D. & Teh, Y. W.) 1799–1808 (PMLR, 2017).
38. Tian, H. & Rangarajan, S. Predicting adsorption energies using multifidelity data. *J. Chem. Theory Comput.* **15**, 5588–5600. <https://doi.org/10.1021/ACS.JCTC.9B00336> (2019).
39. Heyd, J., Scuseria, G. E. & Ernzerhof, M. Hybrid functionals based on a screened Coulomb potential. *J. Chem. Phys.* **118**, 8207. <https://doi.org/10.1063/1.1564060> (2003).
40. Jie, J. S. *et al.* A new MaterialGo database and its comparison with other high-throughput electronic structure databases for their predicted energy band gaps. *Sci. China Technol. Sci.* **62**, 1423–1430. <https://doi.org/10.1007/S11431-019-9514-5> (2019).
41. Sun, J., Ruzsinszky, A. & Perdew, J. P. Strongly constrained and appropriately normed semilocal density functional. *Phys. Rev. Lett.* **115**, 036402. <https://doi.org/10.1103/PhysRevLett.115.036402> (2015).
42. Borlido, P. *et al.* Large-scale benchmark of exchange-correlation functionals for the determination of electronic band gaps of solids. *J. Chem. Theory Comput.* **15**, 5069–5079. <https://doi.org/10.1021/ACS.JCTC.9B00322> (2019).
43. Canning, A., Chaudhry, A., Boutchko, R. & Grønbech-Jensen, N. First-principles study of luminescence in ce-doped inorganic scintillators. *Phys. Rev. B* **83**, 125115. <https://doi.org/10.1103/PhysRevB.83.125115> (2011).
44. Polman, A., Knight, M., Garnett, E. C., Ehrler, B. & Sinke, W. C. Photovoltaic materials: Present efficiencies and future challenges. *Science* <https://doi.org/10.1126/SCIENCE.AAD4424> (2016).
45. Castelli, I. E. *et al.* Computational screening of perovskite metal oxides for optimal solar light capture. *Energy Environ. Sci.* **5**, 5814–5819. <https://doi.org/10.1039/C1EE02717D> (2012).
46. Wu, Y., Lazic, P., Hautier, G., Persson, K. & Ceder, G. First principles high throughput screening of oxynitrides for water-splitting photocatalysts. *Energy Environ. Sci.* **6**, 157–168. <https://doi.org/10.1039/C2EE23482C> (2013).
47. Suram, S. K., Newhouse, P. F. & Gregoire, J. M. High throughput light absorber discovery, part 1: An algorithm for automated tauc analysis. *ACS Combin. Sci.* **18**, 673–681 (2016).
48. Kiselyova, N. N., Dudarev, V. A. & Korzhuyev, M. A. Database on the bandgap of inorganic substances and materials. *Inorg. Mater. Appl. Res.* **7**, 34–39. <https://doi.org/10.1134/S2075113316010093> (2016).
49. Rohr, B. *et al.* Benchmarking the acceleration of materials discovery by sequential learning. *Chem. Sci.* **11**, 2696–2706. <https://doi.org/10.1039/c9sc05999g> (2020).
50. Ward, L. *et al.* Matminer: An open source toolkit for materials data mining. *Comput. Mater. Sci.* **152**, 60–69. <https://doi.org/10.1016/j.commatsci.2018.05.018> (2018).
51. Ong, S. P. *et al.* The materials application programming interface (API): A simple, flexible and efficient API for materials data based on REpresentational State Transfer (REST) principles. *Comput. Mater. Sci.* **97**, 209–215. <https://doi.org/10.1016/J.COMMA TSCI.2014.10.037> (2015).
52. Perdew, J. P., Burke, K. & Ernzerhof, M. Generalized gradient approximation made simple. *Phys. Rev. Lett.* **77**, 3865–3868. <https://doi.org/10.1103/PhysRevLett.77.3865> (1996).
53. Morales-García, Á., Valero, R. & Illas, F. Morphology of TiO₂ nanoparticles as a fingerprint for the transient absorption spectra: Implications for photocatalysis. *J. Phys. Chem. C* **124**, 11819–11824. <https://doi.org/10.1021/ACS.JPCCC.0C02946> (2020).
54. Srinivas, N., Krause, A., Kakade, S. M. & Seeger, M. W. Information-theoretic regret bounds for gaussian process optimization in the bandit setting. *IEEE Trans. Inf. Theory* **58**, 3250–3265. <https://doi.org/10.1109/TIT.2011.2182033> (2012).
55. Desautels, T., Krause, A. & Burdick, J. W. Parallelizing exploration-exploitation tradeoffs in gaussian process bandit optimization. *J. Mach. Learn. Res.* **15**, 4053–4103 (2014).
56. GPy. GPy: A gaussian process framework in python. <http://github.com/SheffieldML/GPy> (since 2012).
57. Hu, S., Xiang, C., Haussener, S., Berger, A. D. & Lewis, N. S. An analysis of the optimal band gaps of light absorbers in integrated tandem photoelectrochemical water-splitting systems. *Energy Environ. Sci.* **6**, 2984–2993. <https://doi.org/10.1039/C3EE40453F> (2013).
58. Belsky, A., Hellenbrandt, M., Karen, V., Luksch, P., IUCr. New developments in the inorganic crystal structure database (ICSD): Accessibility in support of materials research and design. *Acta Crystallogr. Sect. B Struct. Sci.* **58**, 364–369. <https://doi.org/10.1107/S0108768102006948> (2002).
59. Choudhury, R., Aykol, M., Gratzl, S., Montoya, J. & Hummelshøj, J. MaterialNet: A web-based graph explorer for materials science data. *J. Open Source Softw.* **5**, 2105. <https://doi.org/10.21105/joss.02105> (2020).

Acknowledgements

This work was supported by Toyota Research Institute through the Accelerated Materials Design and Discovery program. The authors gratefully acknowledge helpful discussions with Chirranjeevi Gopal, Abraham Anapolsky, and Linda Hung.

Author contributions

A.P. conducted the active learning simulations and compiled the results presented herein. A.P. and J.H.M. prepared and reviewed the code published alongside the manuscript. M.A., S.S., J.H., and J.H.M. conceptualized the project. A.P. and S.B.T. prepared and revised the manuscript figures. All authors wrote and reviewed the manuscript.

Competing interests

The authors declare no competing interests.

Additional information

Supplementary Information The online version contains supplementary material available at <https://doi.org/10.1038/s41598-022-08413-8>.

Correspondence and requests for materials should be addressed to J.H.M.

Reprints and permissions information is available at www.nature.com/reprints.

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



Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

© The Author(s) 2022