ASME Journal of Mechanical Design
Online journal at:
https://asmedigitalcollection.asme.org/mechanicaldesign

Asynchronous Multi-Information Source Bayesian Optimization

Danial Khatamsaz[1]
Department of Materials Science and Engineering,
Texas A&M University,
College Station, TX 77843
e-mail: khatamsaz@tamu.edu

Raymundo Arroyave
Department of Materials Science and Engineering,
Texas A&M University,
College Station, TX 77843
e-mail: rarroyave@tamu.edu

Douglas L. Allaire
J. Mike Walker '66 Department of Mechanical Engineering,
Texas A&M University,
College Station, TX 77843
e-mail: dallaire@tamu.edu

Resource management in engineering design seeks to optimally allocate while maximizing the performance metrics of the final design. Bayesian optimization (BO) is an efficient design framework that judiciously allocates resources through heuristic-based searches, aiming to identify the optimal design region with minimal experiments. Upon recommending a series of experiments or tasks, the framework anticipates their completion to augment its knowledge repository, subsequently guiding its decisions toward the most favorable next steps. However, when confronted with time constraints or other resource challenges, bottlenecks can hinder the traditional BO's ability to assimilate knowledge and allocate resources with efficiency. In this work, we introduce an asynchronous learning framework designed to utilize idle periods between experiments. This model adeptly allocates resources, capitalizing on lower fidelity experiments to gather comprehensive insights about the target objective function. Such an approach ensures that the system progresses uninhibited by the outcomes of prior experiments, as it provisionally relies on anticipated results as stand-ins for actual outcomes. We initiate our exploration by addressing a basic problem, contrasting the efficacy of asynchronous learning against traditional synchronous multi-fidelity BO. We then employ this method to a practical challenge: optimizing a specific mechanical characteristic of a dual-phase steel. [DOI: 10.1115/1.4065064]

Keywords: data-driven design, design of experiments, design optimization, design process, machine learning, simulation-based design

1 Introduction

Recent developments in efficient design and optimization frameworks have prominently featured Bayesian optimization (BO) as a central element. This efficient technique significantly reduces the cost demands tied to optimizing expensive objective functions. BO approaches are capable of working with minimal data and are driven via a heuristic-based search in the design space to discover optimal design regions. Such characteristics have established BO as an exceptionally efficient design technique. It adeptly allocates resources to the most potentially informative experiments, facilitating enhanced learning and yielding superior estimates of the optimal solution.

In BO, a surrogate of the objective function $f(\mathbf{x})$ is constructed, given existing observations, to predict the response across the design domain at unobserved locations. A Gaussian process (GP) is a popular surrogate regression model that is easy to manipulate, for instance, by adding new observations or updating the hyperparameters, as well as being a computationally cheap source for probabilistic prediction of the objective function response [1]. BO heuristic-based search relies on GP probabilistic predictions to evaluate the potential improvement of the system's knowledge using an acquisition function. A proper acquisition function should balance the exploitation of the current system's knowledge and exploration of unobserved design regions. Accordingly, BO iteratively proposes new experiments to learn and discover the optimal regions with minimal resource investment [2,3].

BO establishes an efficient learning framework by allocating resources to the experiment (or, more generally, design point) with a higher chance of improving the knowledge of the system being optimized. However, the characteristics of a particular design problem may drastically increase the need for more observations to make reliable decisions. For instance, increasing the number of design variables requires learning and exploring a vast design space, similarly, confronting an irregular objective function response surface necessitates the need for making more observations to build more accurate GPs. To address the deteriorating performance of BO in such conditions, different BO-based design frameworks have been developed by integrating computational tools to improve modeling and decision-making procedures.

For example, batch BO [4,5], suggests parallel experiments to increase the learning rate within a limited time window. It also eliminates the GP hyperparameter tuning step by considering every possible representation of the objective function. This essentially addresses the issue of working with minimal data in sparse spaces where any systematic hyperparameter tuning approach can fail.

Another example is the use of the active subspace method (ASM) [6–8] to define the design problem on low-dimensional spaces to speed up the learning process. ASM-based BO focuses on subspaces with larger objective function variability to maximize potential improvements of a target value in a fewer number of iterations by ignoring less informative regions of the design space.

[1]Corresponding author.

Contributed by Design Automation Committee of ASME for publication in the Journal of Mechanical Design. Manuscript received October 16, 2023; final manuscript received March 8, 2024; published online April 9, 2024. Assoc. Editor: Xueguan Song.

To significantly decrease the number of expensive experiments in Bayesian design campaigns, multi-fidelity BO frameworks [9–13] have been proposed to exploit information from multiple sources instead of relying on one expensive model or experiment. In many engineering applications, there exists more than one experiment/model, which we call an information source, to estimate a target property, each with a different fidelity level and evaluation cost. It is assumed that any information source has potentially useful information about the target property, thus, BO leverages resources to exploit multiple information sources. Then, a fused model is constructed by fusing information from all contributing sources. The information is perceived by correlating the responses of lower fidelity models to each other and to the highest fidelity source, known as the ground truth. The queries from the ground truth model are driven via information learned from low-fidelity sources to get the most out of the limited allowed number of experiments.

Physics-informed BO is another proposed methodology to augment knowledge from underlying governing laws of physical phenomenons or partially observable behavior of systems to lower the dependency on statistical information in the form of input–output data [14–17].

The aforementioned BO-based frameworks have been shown to outperform conventional BO by allocating resources effectively to more valuable experiments to collect information whether from multiple cheaper sources or a single source. However, a common limitation of any BO-based design framework is that the design process comes to a halt while waiting for completion of the required experiments. Once the BO suggests the next experiment to run, it has to wait until the result is returned to make the next decision. This so-called synchronous learning prevents the framework from optimally learning the objective function under a time constraint.

Some works have been done to propose asynchronous learning scenarios to get over the halting challenge. One suggested approach to handling multiple independent experiments [18,19] is to do Thompson sampling to ensure variety in experiments. Additionally, penalizing already sampled regions is proposed to force the system to explore non-sampled or less-sampled regions [20] while previously selected experiments are still running. Another technique is to sample from GP posterior to estimate an experiment's outcome, then, temporarily update the GP and compute possible acquisition function values to make the next decision [21]. This, in fact, follows a Monte Carlo sampling approach to generate a sufficiently large number of samples from a yet-to-be-completed experiment posterior to compute the expected value of the acquisition function. However, this latter approach can get computationally intractable if multiple sequential decisions should be made.

In Ref. [22], an asynchronous framework is suggested that distributes experiments among multiple workers (cores) and as soon as one core is released, the system uses the result to make the next decision. Thus, the process is not "bottlenecked" by the slowest experiment. However, the downside of this approach is that assuming a batch of m experiments are running, the next decision is made while no information is available on $m-1$ experiments. Although the system is moving forward and proposing asynchronous experiments, it may not be the most efficient approach to decision-making.

In this work, we propose an asynchronous multi-fidelity BO with a decision-making strategy to partially incorporate information from running experiments *yet to be completed*. This is done by triggering an experiment and temporarily augmenting the optimistic expected outcome, based on the expected improvement computation, to update the system's knowledge to move forward while the experiment is still running.

Accounting for expected improvement is crucial. Augmenting the mean prediction of a GP, as the expected outcome, at a location $\mathbf{x}$ to itself does not change the mean response surface of the GP since it is essentially passing through the predicted value. The immediate impact, however, is shrinkage of the uncertainty around $\mathbf{x}$ as it is assumed to be the new observed data with no uncertainty. The reduction in uncertainty results in potentially smaller expected knowledge gain evaluated by the acquisition function in the neighborhood of $\mathbf{x}$ and BO ignorance to explore there. Thus, not much information is added to the system to impact the next decision other than suggesting a different experiment—which could be done by batch BO—without interaction among the selected experiments. Our proposed approach also enables us to perform multi-fidelity BO to select both design and information source to query since each information source could be updated independently. This is also computationally cheaper in comparison to the Monte Carlo sampling method to compute the expected value of an acquisition function while multiple experiments are still running.

In the following, we first introduce the ingredients of our multi-fidelity BO framework such as GPs, information fusion, and the search policy. Then, the main steps taken to complete each iteration of the framework are discussed. In Sec. 3, we use a synthetic function to show how the framework exploits multiple information sources in both synchronous and asynchronous scenarios. Then, we show the application in designing a dual-phase steel. In Sec. 4, we summarize the work and discuss potential improvements to the framework to address more realistic applications.

2 Methods

We consider an optimization problem as

$$\mathbf{x}^* = \arg\max_{\mathbf{x}\in\chi} f(\mathbf{x}) \tag{1}$$

where $f(\mathbf{x})$ is an expensive-to-evaluate black-box objective function defined on the design space χ. The goal of the optimization process is to approximate $\mathbf{x}^*$ as accurately as possible by spending limited resources assigned to the task. We assume that more than one computational model or experiment exists to approximate the objective function at different design points. Thus, we seek a multi-fidelity approach, suggesting a more efficient optimization scheme to locate the optimum design region by lowering the number of expensive function evaluations. By using a multi-fidelity setting, we build a predictive model to approximate the response of the expensive black-box objective function by fusing information from multiple models/experiments, known as information sources. In the following, we discuss different ingredients of the optimization framework and then, the approach to learn from multiple information sources asynchronously.

The first step to employing a multi-fidelity Bayesian optimization framework is constructing surrogates of each information source to connect the design space to each source's response. A Gaussian process is commonly used as the surrogate to not only predict the response but also provide the uncertainty associated with its predictions [1]. A Gaussian process conditions a probabilistic function to the training data and provides a normal distribution as its prediction at an unobserved location.

In the presence of i information sources and N_i previously evaluated points: $\{\mathbf{X}_{N_i}, \mathbf{y}_{N_i}\}$, where $\mathbf{X}_{N_i} = (\mathbf{x}_{1,i}, \ldots, \mathbf{x}_{N_i,i})$ are N_i input samples and $\mathbf{y}_{N_i} = \big(f_i(\mathbf{x}_{1,i}), \ldots, f_i(\mathbf{x}_{N_i,i})\big)$ are the respective objective values, information source i modeled by a Gaussian process at input location $\mathbf{x}$ is represented by

$$f_{\mathrm{GP},i}(\mathbf{x}) \mid \mathbf{X}_{N_i}, \mathbf{y}_{N_i} \sim \mathcal{N}\big(\mu_i(\mathbf{x}), \sigma^2_{\mathrm{GP},i}(\mathbf{x})\big) \tag{2}$$

where

$$\begin{aligned}\mu_i(\mathbf{x}) &= K_i(\mathbf{X}_{N_i}, \mathbf{x})^T [K_i(\mathbf{X}_{N_i}, \mathbf{X}_{N_i}) + \sigma^2_{n,i} I]^{-1} \mathbf{y}_{N_i} \\ \sigma^2_{\mathrm{GP},i}(\mathbf{x}) &= k_i(\mathbf{x}, \mathbf{x}) - K_i(\mathbf{X}_{N_i}, \mathbf{x})^T \\ &\quad [K_i(\mathbf{X}_{N_i}, \mathbf{X}_{N_i}) + \sigma^2_{n,i} I]^{-1} K_i(\mathbf{X}_{N_i}, \mathbf{x})\end{aligned} \tag{3}$$

where k_i is a real-valued kernel function over the input space, $K_i(\mathbf{X}_{N_i}, \mathbf{X}_{N_i})$ is the $N_i \times N_i$ matrix with m, n entry as $k_i(\mathbf{x}_{m,i}, \mathbf{x}_{n,i})$, and $K_i(\mathbf{X}_{N_i}, \mathbf{x})$ is the $N_i \times 1$ vector with mth entry as $k_i(\mathbf{x}_{m,i}, \mathbf{x})$. The term $\sigma^2_{n,i}$ models observation error that stems from experiments.

A common choice for kernel is the squared exponential covariance function

$$k_i(\mathbf{x}, \mathbf{x}') = \sigma_s^2 \exp\left(-\sum_{h=1}^{d} \frac{(x_h - x_h')^2}{2l_h^2}\right) \tag{4}$$

where d indicates the dimensionality of the input space, σ_s^2 is the signal variance, and l_h, where $h = 1, 2, \ldots, d$, is the characteristic length-scale that defines the correlation strength in each dimension.

The next step in constructing a multi-fidelity framework is to employ information fusion technique to integrate information from all sources into a single predictive model. This way, we are able to spend resources on evaluating different information sources such that the predictive model encapsulates the highest amount of information about the optimum design at each stage of the optimization process [23–25].

There are different information fusion techniques such as Bayesian modeling averaging [26–31], the use of adjustment factors [32–35], covariance intersection methods [36,37], and fusion under known correlation [38–40]. In our multi-fidelity optimization framework, the optimization is performed with respect to the ground truth regardless of the hierarchy of fidelity. This technique uses reification to calculate the correlation between pairs of information sources to fuse their information and approximate the ground truth [23,25].

Reification has been proposed in Ref. [40] as a technique to combine (fuse) multiple normal probability distributions as dependent information sources. Here, since information sources are represented via Gaussian processes, "Reification" is an appropriate approach to fuse multiple information sources assuming they are all correlated as they are approximating the same objective function. Following Ref. [40], the fused mean and fused variance estimates from multiple predictions in the forms of normal distributions are defined as

$$\mathbb{E}[\hat{f}(\mathbf{x})] = \frac{\mathbf{e}^{\mathrm{T}}\tilde{\Sigma}(\mathbf{x})^{-1}\boldsymbol{\mu}(\mathbf{x})}{\mathbf{e}^{\mathrm{T}}\tilde{\Sigma}(\mathbf{x})^{-1}\mathbf{e}} \tag{5}$$

$$\mathrm{Var}\left(\hat{f}(\mathbf{x})\right) = \frac{1}{\mathbf{e}^{\mathrm{T}}\tilde{\Sigma}(\mathbf{x})^{-1}\mathbf{e}} \tag{6}$$

where $\mathbf{e} = [1, \ldots, 1]^{\mathrm{T}}$, $\boldsymbol{\mu}(\mathbf{x}) = [\mu_1(\mathbf{x}), \ldots, \mu_m(\mathbf{x})]^{\mathrm{T}}$ given m information sources estimating the objective value at input location $\mathbf{x}$, and $\tilde{\Sigma}(\mathbf{x})$ is the covariance matrix between information sources

$$\tilde{\Sigma}(\mathbf{x}) = \begin{bmatrix} \sigma_1^2(\mathbf{x}) & \cdots & \rho_{1m}(\mathbf{x})\sigma_1(\mathbf{x})\sigma_m(\mathbf{x}) \\ \rho_{12}(\mathbf{x})\sigma_1(\mathbf{x})\sigma_2(\mathbf{x}) & \cdots & \rho_{2m}(\mathbf{x})\sigma_2(\mathbf{x})\sigma_m(\mathbf{x}) \\ \vdots & \ddots & \vdots \\ \rho_{1m}(\mathbf{x})\sigma_1(\mathbf{x})\sigma_m(\mathbf{x}) & \cdots & \sigma_m^2 \end{bmatrix} \tag{7}$$

where $\rho_{i,j}$ is the correlation coefficient between the deviations of information sources i and j at $\mathbf{x}$. To estimate the correlation coefficient, information sources are reified in turn which means they are treated as ground truth. Then the correlation between deviations is calculated as

$$\rho_{ij}(\mathbf{x}) = \frac{\sigma_j^2(\mathbf{x})}{\sigma_i^2(\mathbf{x}) + \sigma_j^2(\mathbf{x})}\bar{\rho}_{ij}(\mathbf{x}) + \frac{\sigma_i^2(\mathbf{x})}{\sigma_i^2(\mathbf{x}) + \sigma_j^2(\mathbf{x})}\bar{\rho}_{ji}(\mathbf{x}) \tag{8}$$

where $\bar{\rho}_{ij}$ and $\bar{\rho}_{ji}$ are

$$\bar{\rho}_{ij}(\mathbf{x}) = \frac{\sigma_i(\mathbf{x})}{\sqrt{\left(\mu_i(\mathbf{x}) - \mu_j(\mathbf{x})\right)^2 + \sigma_i^2(\mathbf{x})}} \tag{9}$$

$$\bar{\rho}_{ji}(\mathbf{x}) = \frac{\sigma_j(\mathbf{x})}{\sqrt{\left(\mu_i(\mathbf{x}) - \mu_j(\mathbf{x})\right)^2 + \sigma_j^2(\mathbf{x})}} \tag{10}$$

The root squared term in the denominators in Eqs. (9) and (10) are representing the total uncertainty of each information source as the sum of the information source variance and square of the deviation with respect to the reified information source. The information source variance itself is defined as

$$\sigma_i^2(\mathbf{x}) = \sigma_{\mathrm{GP},i}^2(\mathbf{x}) + \sigma_{\mathrm{disc},i}^2(\mathbf{x}) \tag{11}$$

that is the total uncertainty that stems from the GP probabilistic modeling and the discrepancy with respect to the ground truth model. A higher fidelity information source has less discrepancy with respect to the ground truth. Consequently, it is less correlated to lower fidelity sources. In general, the reification process puts larger weights on less uncertain information sources when calculating the linear combination of the distributions.

The fused model, also a Gaussian process, is then created using fused means and fused variances as

$$\hat{f}^{\mathrm{fused}}(\mathbf{x}) \sim \mathcal{N}(\mu^{\mathrm{fused}}(\mathbf{x}), \Sigma^{\mathrm{fused}}(\mathbf{x})) \tag{12}$$

The heuristic we employ in the context of BO should address the trade-off between the exploration of unobserved locations and the exploitation of the current system's knowledge. In a multi-fidelity setting, the decision should be made on what design and which information source to query to gain the most information out of an experiment. Here, we use the knowledge gradient (KG) policy [41–43] as the metric that quantifies the expected change in the system's knowledge of the optimum quantity of interest when evaluating a potential design point using a given "information source." To evaluate the expected knowledge gain if a design point $\mathbf{x}$ is added, we define information sources, already queried design points, and corresponding objective values by $(i_{1:N}, \mathbf{x}_{1:N}, y_{1:N})$ for the first N queries. $\hat{f}$ is the posterior distribution of the fused model and the expected improvement is given as

$$\begin{aligned} \mathrm{EI}(\mathbf{x}) &= \mathbb{E}\left[\max_{\mathbf{x}' \in \mathcal{X}} \mathbb{E}[\hat{f}(\mathbf{x}')|i_{1:N}, \mathbf{x}_{1:N}, \mathbf{x}_{N+1} = \mathbf{x}, y_{1:N}] - \max_{\mathbf{x}' \in \mathcal{X}} \mathbb{E}[\hat{f}(\mathbf{x}')|i_{1:N}, \mathbf{x}_{1:N}, y_{1:N}]\right] \\ &= \mathbb{E}\left[\max_{\mathbf{x}' \in \mathcal{X}} \mathbb{E}[\hat{f}(\mathbf{x}')|i_{1:N}, \mathbf{x}_{1:N}, \mathbf{x}_{N+1} = \mathbf{x}, y_{1:N}]\right] - \max_{\mathbf{x}' \in \mathcal{X}} \mathbb{E}[\hat{f}(\mathbf{x}')|i_{1:N}, \mathbf{x}_{1:N}, y_{1:N}] \end{aligned} \tag{13}$$

Note that the last expression is known and can be removed from the expectation operator as it is conditioned on the first N queries. Then, using KG policy to maximize this expectation, the value of being at the knowledge state H^N is defined as $V^N(H^N) = \max_{\mathbf{x} \in \mathcal{X}} H^N$, where the knowledge state itself is presented by $H^N = \mathbb{E}[\hat{f}(\mathbf{x}')|i_{1:N}, \mathbf{x}_{1:N}, y_{1:N}]$. The KG as a measure of expected improvement is written as

$$\nu^{\mathrm{KG}}(\mathbf{x}) = \mathbb{E}[V^{N+1}(H^{N+1}(\mathbf{x})) - V^N(H^N)|H^N] \tag{14}$$

The decision on the information source to query is made by repeating the calculations, each time updating a different information source to select the best pair of a design point and an information source that maximizes $\nu^{KG}(\mathbf{x})$ as the next BO experiment.

A multi-fidelity BO framework is created by implementing Gaussian processes as surrogates, an information fusion technique,

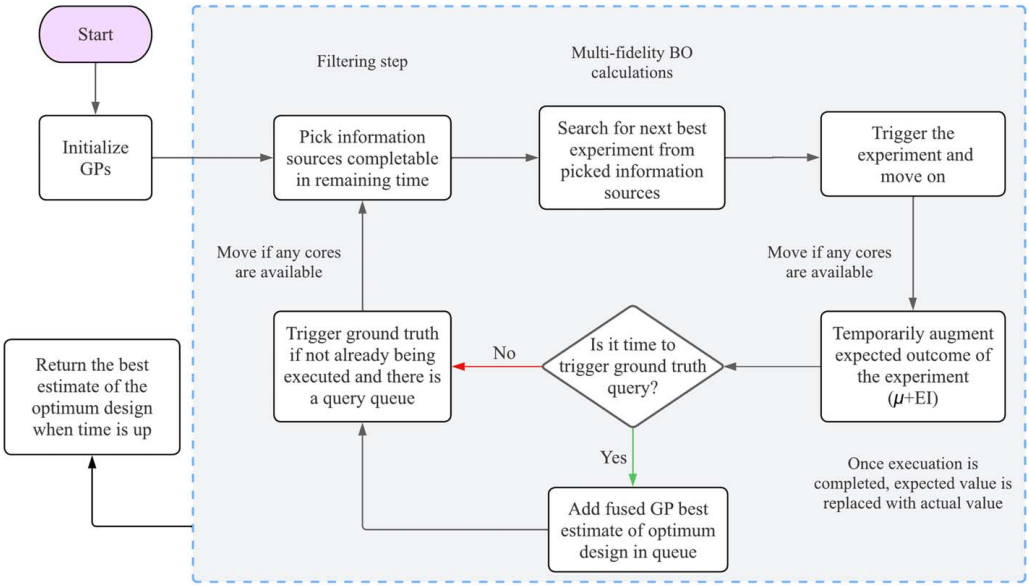

Fig. 1 Schematic of the proposed asynchronous multi-information source Bayesian optimization. The main loop starts by filtering information sources, then, making the decision on design and information source to query. Once the experiment is triggered, the result is temporarily augmented until the true result is returned. The ground truth is queried on a timely basis to update information source discrepancy terms.

which is reification here, and an acquisition function such as KG. Then heuristic-based decisions are made based on the constructed fused model's predictions and acquisition function evaluations on potential next experiments to select the most informative experiments to discover the optimum design [5,8,10,11,13,24,44–46].

While it has been shown that taking a multi-fidelity approach increases the design efficiency significantly, it may not use available resources as efficiently considering there is idle time while waiting for the completion of an experiment. More specifically, once a decision is made about the next experiment to run, an information source will be called and the BO loop will be paused, awaiting the experiment to be completed.

Here, we propose an approach to exploit resources more efficiently by preventing a halt in the optimization process. In this approach, an experiment is triggered, but the optimistic expected outcome of the experiment is taken as a temporary result as long as the experiment is running. Then, the optimization loop moves on to make the next decision based on the temporary updated system knowledge. This leads to asynchronous calculations while the process is iterating forward.

To calculate the optimistic expected outcome of an experiment, we use expected improvement, $\mathrm{EI}(\mathbf{x})$, as a systematic prediction of the expected improvement at $\mathbf{x}$. Then the optimistic expected outcome is $\mu(\mathbf{x}) + \mathrm{EI}(\mathbf{x})$ where $\mu(\mathbf{x})$ is the mean of the GP posterior distribution at design point $\mathbf{x}$. The expected improvement formula, assuming the GP posterior distribution defined by $\mathcal{N}(\mu(\mathbf{x}), \sigma^2(\mathbf{x}))$, is calculated as

$$\begin{aligned}\mathrm{EI}(\mathbf{x}) &= \mathbb{E}[\max(y, f_{\text{best}}) - f_{\text{best}}] \\ &= (\mu(\mathbf{x}) - f_{\text{best}})\Phi\left(\frac{(\mu(\mathbf{x}) - f_{\text{best}})}{\sigma(\mathbf{x})}\right) + \sigma(\mathbf{x})\phi\left(\frac{(\mu(\mathbf{x}) - f_{\text{best}})}{\sigma(\mathbf{x})}\right)\end{aligned} \tag{15}$$

where Φ and ϕ are cumulative distribution function (CDF) and probability density function (PDF) of the standard normal distribution respectively. However, we are using expected improvement as a systematic way to define how optimistic the framework should be, thus, the expected improvement is calculated with respect to $\mu(\mathbf{x})$ and Eq. (15) reduces down to approximately $0.4 \times \sigma(\mathbf{x})$. In the synchronous scenario, a GP is updated once the experiment is completed by augmenting the result as a new training data $\{\mathbf{x}, f(\mathbf{x})\}$, however, in the asynchronous scenario, a GP is updated immediately after decision-making using the expected outcome, $\{\mathbf{x}, \mu(\mathbf{x}) + \mathrm{EI}(\mathbf{x})\}$. Once the experiment is completed, the augmented data point in the GP will be updated with the actual experimental value.

Here, we assume that the budget for optimization is given as a specific number of cores (or compute resources) under a time constraint. We point out that this assumption is made to motivate the case study. However, the same framework can potentially be used in a non-computational setting as long as there are hard resource constraints (e.g., time) and there are multiple tasks that one can carry out asynchronously.

Figure 1 illustrates the steps in our proposed approach. The overall goal here is to allocate resources as efficiently as possible to the optimization process. Each step is explained in more detail in the following:

(1) The system picks information sources that, if selected, can be completed in the remaining time window.
(2) Multi-fidelity BO calculations are performed by creating a fused model from temporary updated information sources. The experiment that maximizes the acquisition function will be selected.
(3) A selected experiment is triggered and if any cores are left, it moves forward.
(4) While the experiment is running, the optimistic expected outcome, $\mu(\mathbf{x}) + \mathrm{EI}(\mathbf{x})$, is temporarily augmented to the respective GP. Once the experiment is completed, the experimental value replaces the augmented expected value.
(5) The ground truth model is queried on a timely basis to update system knowledge regarding information source deviations from the true model. Since only one ground truth calculation is allowed at any time, a queue of queries is created to ensure queries are done in order.

With this approach, if multiple cores are given to a design task, it is ensured that all cores carry out calculations at all times and the resources are fully exploited to increase the learning rate of the Bayesian design campaign. The only situation in the optimization, when the iterative loop comes to a halt is when all cores are being used to run experiments. However, as soon as one experiment

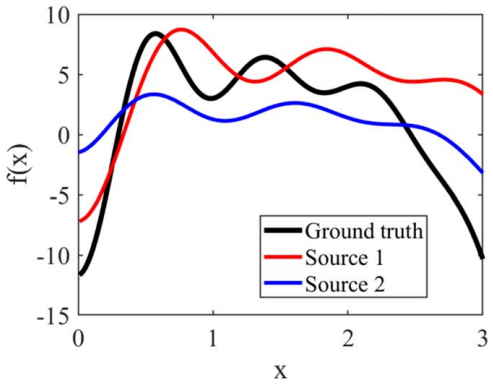

Fig. 2 Synthetic functions representing a ground truth model and two lower fidelity sources. The goal is to find the ground truth's maximizer.

is completed, the core is released and carries the rest of the calculations. Once the time is up, the framework returns the best estimate of the optimum design predicted by the fused model.

Note that with this framework, experiments are triggered regardless of other experiments' status, which leads to asynchronous learning from multiple sources. To program this framework, we used asynchronous functions in MATLAB to create future objects that store the status of the experiment and the final results when completed. Thus, the status of a future object can be monitored to read the final result as soon as the calculations are finished.

3 Results

To investigate the performance of our proposed asynchronous active learning framework, first, we solve two synthetic test problems and then we apply it to a real-world engineering application to optimize the mechanical response of dual-phase steel. In all test problems, the information source GPs are initialized with one random training data and we assumed there is no observation from the ground truth yet at the beginning of the optimization process. In all cases, the results will be compared to the synchronous version of our multi-fidelity Bayesian optimization framework.

The maximization test problem is designed by creating a ground truth model and two information sources to approximate the response of the ground truth model as shown in Fig. 2. We assume that running each ground truth query requires 10 s and low-fidelity sources each require 2 s to be completed. For this problem, we assume the budget is given as four cores for 100 s to return the best estimate of the optimum design. This is equivalent to 400 s of calculations. The way resources are distributed between a number of cores and wall-time is a subject of future study, however, here, to avoid the complexity caused by varying the number of cores and time window, we stick to four cores and 100 s of computation for each core.

Note that in both the asynchronous and synchronous versions, the same amount of computational resources is allocated. The asynchronous version uses cores to run more experiments to increase the learning rate, however, the synchronous version can only use cores to do parallel computations of acquisition functions in the decision-making step. In both versions, the ground truth query is

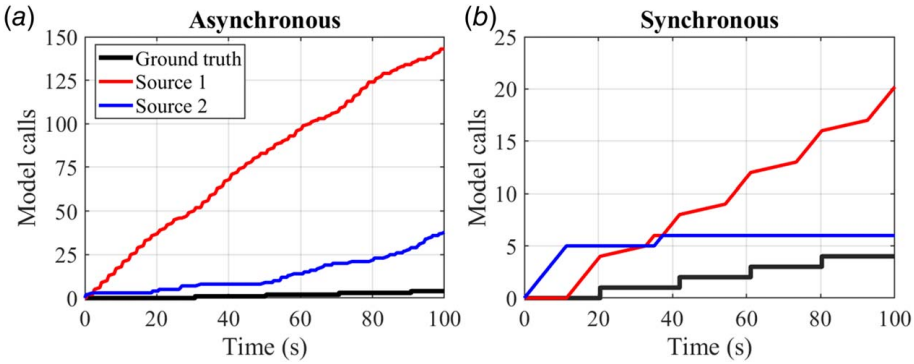

Fig. 3 Exploitation of information sources in (*a*) asynchronous and (*b*) synchronous scenarios. Asynchronous learning enables more queries from information sources by recovering the idle times between experiments and increasing the system's knowledge at a faster pace.

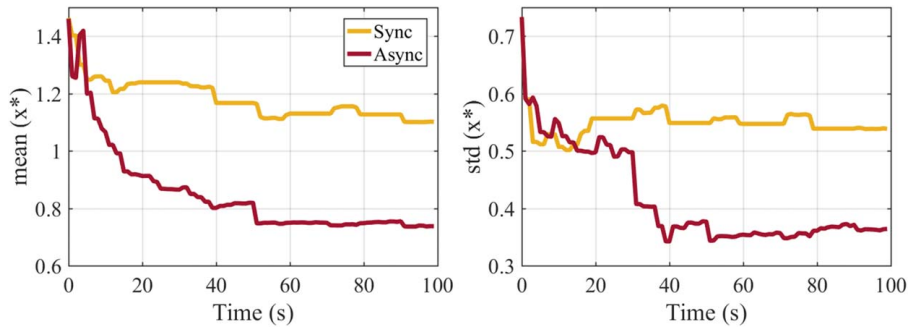

Fig. 4 The mean and standard deviation of the estimated optimum design over 100 replications of simulations as a function of time. More information source queries in asynchronous scenario enhances the system's learning capability and leads to faster convergence to the optimum design region.

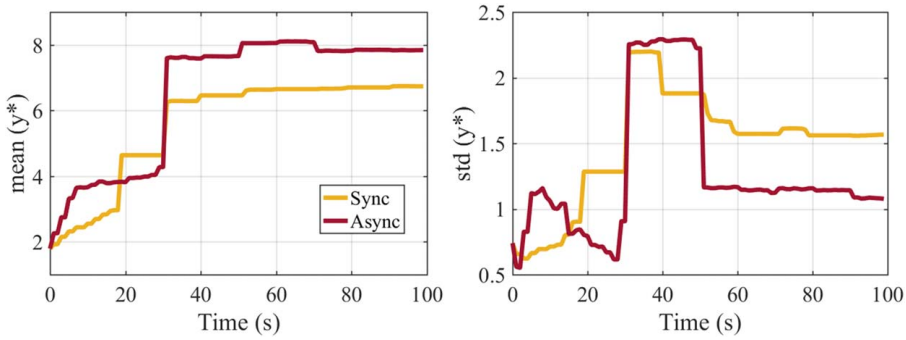

Fig. 5 The fused mean and standard deviation of the discovered optimum objective function value over 100 replications of simulations. The jumps in the left panel are attributed to the initial exploration of the space by the Bayesian optimization framework.

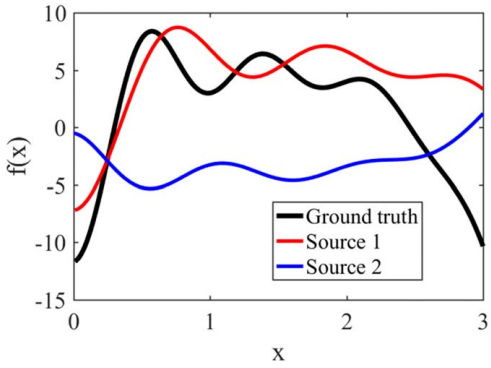

Fig. 6 Synthetic functions representing a ground truth model, a medium-fidelity source, and a low-fidelity source. The goal is to find the ground truth's maximizer

initiated every 20 s to ensure the expensive model is queried on the same basis for fair comparison.

The number of queries from low-fidelity sources and the ground truth model are shown in Fig. 3 for both asynchronous and synchronous versions. As shown in the figure, the asynchronous version makes a larger number of calls to the low-fidelity sources and fully exploits the computational resources while the synchronous version comes to a halt every time an experiment is initiated.

To obtain the average performance of both asynchronous and synchronous versions, the estimated optimum design and corresponding objective function value as a function of time are averaged over 100 replications of simulations, as shown in Figs. 4 and 5, respectively.

On average, asynchronous optimization rapidly finds the optimum design region since the system's knowledge is increasing at a faster pace. This, in fact, shows that if the system exploits other cores' (or other resources) idle time to actively learn by making more observations from information sources, it enhances the overall performance of the system. Note that the synchronous version makes all decisions based on actual experimental values at all times, but the asynchronous version relies on the estimated outcomes of the experiments and continues collecting information. However, making more observations overcompensates the inaccuracy that stems from using expected results.

The decision-making in asynchronous optimization relies on the uncertainties associated with each information source. Therefore, there is a possibility that in initial iterations, when there are not enough ground truth observations, the decision-making is adversely influenced by inaccurate discrepancy calculations. The question here is: what if an information source is very inaccurate and initially, there is not enough evidence to determine such discrepancy to the ground truth? While this is usually the scenario in any multi-fidelity BO process, it may deteriorate the asynchronous optimization performance due to the initial reliance on all sources for a longer time window and the accumulation of GP errors. In such situations, the rationale behind decision-making in the early stages would be the predicted objective function value regardless of its credibility. For instance, in the case of maximizing the objective function, the framework relies on information sources that suggest a larger objective function. Eventually, it learns the discrepancy of sources to the ground truth and decides to discard or keep querying a source. To confirm the capability of the asynchronous scenario to handle such situations, another test problem is defined in Fig. 6 where the second source has a large discrepancy to the ground truth while also showing an opposite trend in objective function variability. We use similar design settings as the previous test

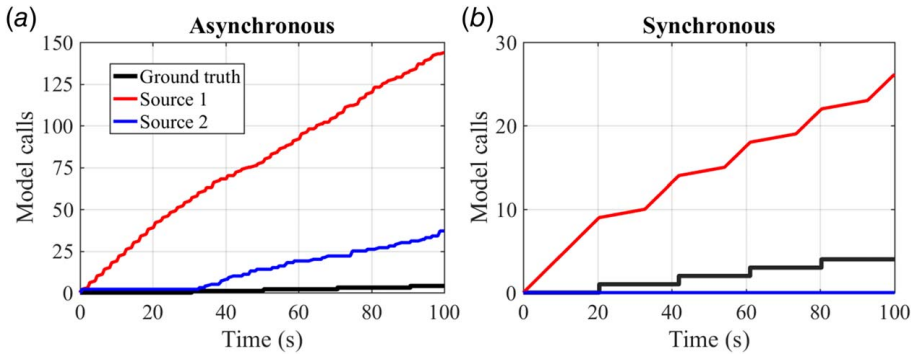

Fig. 7 An instance of exploitation of information sources in (*a*) asynchronous and (*b*) synchronous scenarios

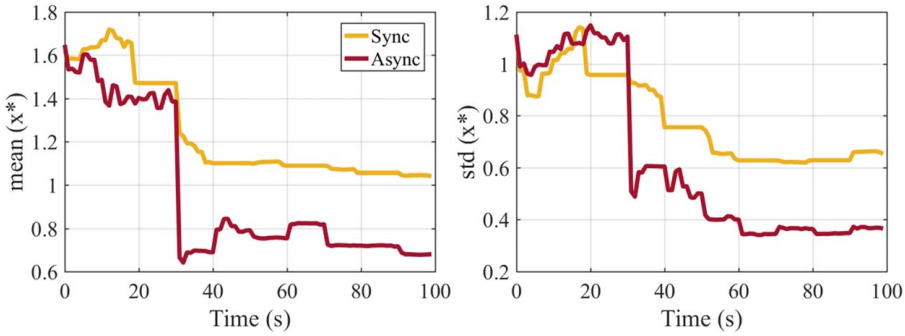

Fig. 8 The mean and standard deviation of the estimated optimum design over 100 replications of simulations as a function of time. More information source queries in asynchronous scenario enhances the system's learning capability and leads to faster convergence to the optimum design region.

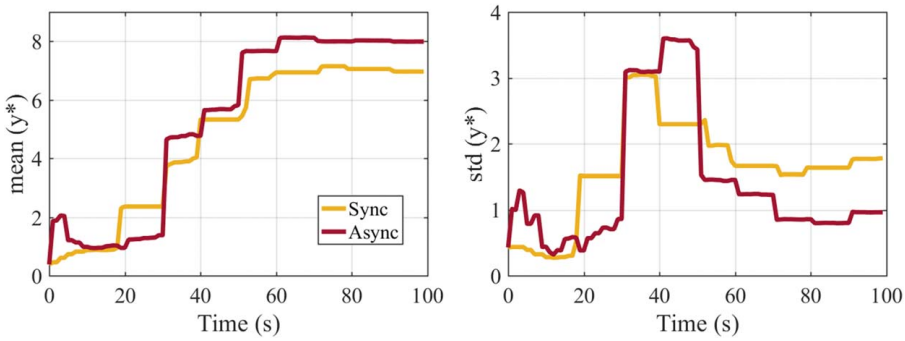

Fig. 9 The fused mean and standard deviation of the discovered optimum objective function value over 100 replications of simulations. The jumps in the left panel are attributed to the initial exploration of the space by the Bayesian optimization framework.

problem. Here, the low-fidelity source also predicts lower objective function values in comparison to the other source, thus, it is expected that the framework initially discards the low-fidelity source.

By looking at the calls to each source in Fig. 7, it is observed that the framework is now calling the lower fidelity source less as it finds no value in collecting information from it. However, note that the low-fidelity source discrepancy to the ground truth is smaller on the right-end side of the design space, and the framework still calls it to collect some useful information.

Figures 8 and 9 also illustrate the optimum design point, the maximum fused objective function, and the standard deviations as functions of time. In contrast to the first test problem, the asynchronous BO here needs more time to recognize and move toward the optimum design region as shown in Fig. 8 since it is collecting information only from one source. After exhausting the first source, it starts to collect information from the second source as well. This is not the case in synchronous optimization as 100 s is probably not enough to exhaust the better source.

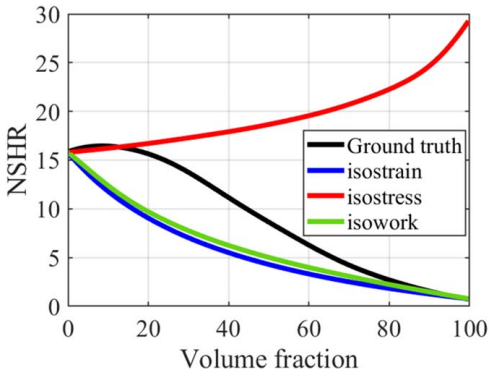

Fig. 10 Normalized strain hardening rates as a function of martensite phase for the ground truth (finite element model) and simplified lower fidelity models: isostrain, isostress, and isowork

We now demonstrate the results of applying our asynchronous active learning framework to maximize the normalized strain hardening rate (NSHR), $(1/\tau)(d\tau/d\varepsilon^{pl})$ at a plastic strain level of 1.5%, of a dual-phase (ferrite-martensite) steel by adjusting the volume fraction of martensite phase—dual-phase steels, a variation of advanced high-strength steels, have experienced fast growth in the automotive industry due to desired microstructural properties [47]. Optimizing their properties requires careful optimization of their microstructure through chemical or processing tuning.

In dual-phase steels, both ferrite and martensite phases undergo nonlinear elastic-plastic deformation with significantly different strength levels and strain hardenability [48,49]. The goal here is to discover the optimum volume fraction of the martensite phase that maximizes NSHR at a given strain level.

To model the mechanical behavior of the dual-phase steel, an expensive finite element model-based simulation is carried out to compute the mechanical response of a three-dimensional representative volume elements of the dual-phase steel microstructure—for more details on this problem, please consult Ref. [44]. However, since the calculations may take hours, in this work, we replace the finite element model with a surrogate (a Gaussian process) to accelerate the design process. On the other hand, there are three simplified lower fidelity models to estimate the mechanical property of interest at lower costs based on isostrain [50], isostress [51], and isowork [52] conditions. These models are discussed in more

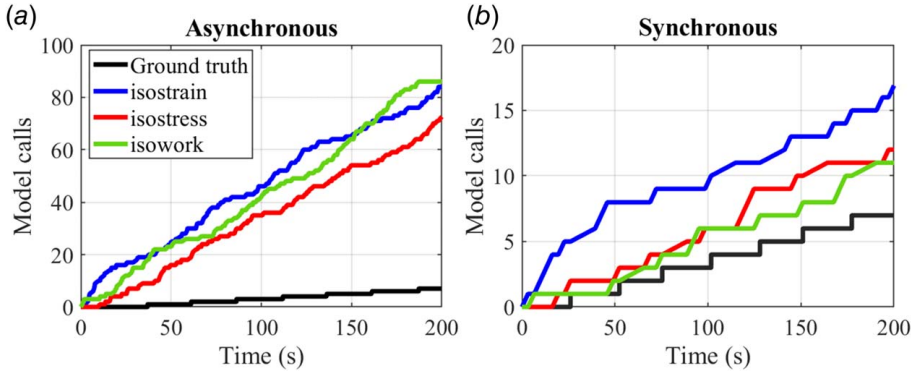

Fig. 11 Exploitation of information sources in (*a*) asynchronous and (*b*) synchronous scenarios. The ground truth model is called on a timely basis in both scenarios.

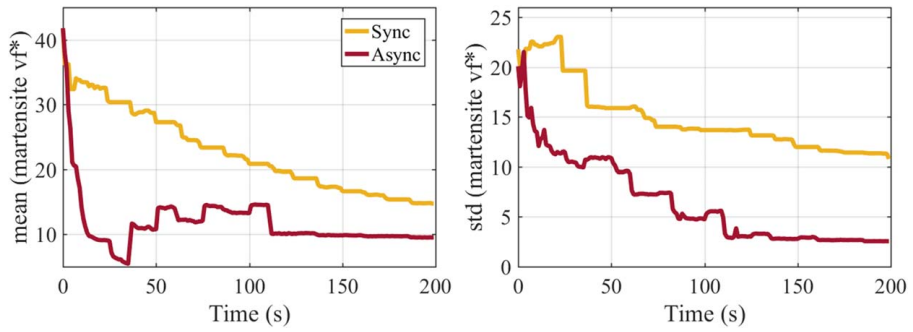

Fig. 12 The mean and standard deviation of the estimated optimum martensite volume fraction over 100 replications of simulations as a function of time. The asynchronous scenario quickly discovers the optimal region and the jumps are caused by replacing the temporary estimations with the actual experiment results.

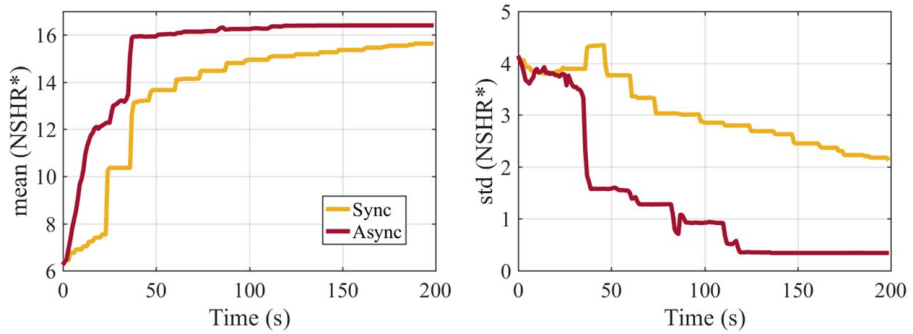

Fig. 13 The fused mean and standard deviation of the discovered optimum normalized strain hardening rate over 100 replications of simulations. Using asynchronous learning technique, the optimum objective function is discovered rapidly.

detail in Refs. [10,45]. Figure 10 shows the response of all models, the normalized strain hardening rate as a function of martensite volume fraction.

For the demonstration problem, the computational resource is given as four cores for 200 s. Each low-fidelity query takes ~3 s and every ground truth query takes 10 s to respond. A ground truth query is initiated every 25 s to correct the discrepancies between information sources and the ground truth model for more accurate correlation calculations between models. The number of model calls is plotted versus time for both asynchronous and synchronous versions in Fig. 11. As shown in the figure, the asynchronous learning framework is capable of recovering the idle times between experiments to make a significantly larger number of observations from each source.

The estimated optimum volume fraction and resulting NSHR values versus time are averaged over 100 replications of simulations and are plotted in Figs. 12 and 13 respectively. Similar to the toy problem's results, the asynchronous version rapidly recognizes the optimal volume fraction range and searches to find the best value. The small jumps are attributed to the fact that the asynchronous framework makes decisions based on temporarily updated knowledge with expected experiment outcomes. It is eventually corrected as more observations are made from information sources and the ground truth. Importantly, as in the case of the toy problem, in this case the asynchronous-BO framework produced results that were significantly less uncertain. This lower uncertainty is an indication of the robustness of the approach and is also evidence that the asynchronous-BO framework effectively

learns more about the problem being solved under the same resource budget.

Note that our proposed asynchronous learning differs from traditional parallel experimentation in the sense that, first, an expected value of previous experiments, even though not completed yet, is incorporated in the decision-making process, while parallel experiments are only decided based on completed experiments. Second, our framework is not bottlenecked by the slowest experiment as is the case in conventional multi-information source BO frameworks. As soon as one experiment is completed, regardless of the other experiments' status, a new decision is made to trigger a new experiment. Also note that similar to batch BO, our sequential asynchronous version also uses multiple cores to handle multiple calculations, but any decision has some degree of dependency on the previously completed or still running experiment(s). In contrast, in batch BO, there is no such interaction between decisions of the same batch of experiments.

4 Conclusions

In this work, we proposed an asynchronous multi-fidelity Bayesian optimization framework to recover the idle times between experiments to increase the learning capacity of the system. In contrast to the synchronous learning framework, with the asynchronous version, the system does not come to a halt waiting for any experiment completion, rather, it calculates the expected outcome of the experiment based on a Gaussian process probabilistic prediction and moves on to make the next decision. In this study, we assumed that the computational resources are given as a number of cores for a limited time window since all the information sources are computational models. In more realistic examples, an information source can be an experimental procedure to measure a property of interest. Then, taking the asynchronous approach suggests initiating the experiment and moving on, however, other types of resource constraints may apply such as financial considerations.

Future work consists of two main aspects of asynchronous resource allocation: distribution of (computational) resources, and decision-making based on additional resource constraints such as financial considerations. The former addresses the question of how many cores should be exploited concerning the number of information sources to maximize the learning rate and the latter takes into account other resource allocation options in addition to time to ensure the most efficient decisions are made considering the respective costs. Additionally, more thoughtful decisions on the selection of information sources may improve the performance of the framework. In this study, the basis for the decision-making process was solely maximization of the knowledge gradient policy, however, the expected gains could be calculated more intelligently by considering the improvements per unit cost of running the experiment or per unit time according to the specific constraints of a design task.

Acknowledgment

Research was sponsored by the Army Research Laboratory and was accomplished under Cooperative Agreement No. W911NF-22-2-0106. The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the Army Research Laboratory or the US Government. The US Government is authorized to reproduce and distribute reprints for Government Purposes notwithstanding any copyright notation herein. We also acknowledge the support of the US Department of Energy (DOE) ARPA-E ULTIMATE Program through Project DE-AR0001427. D.K. acknowledges the support of NSF through Grant No. NSF-DMREF-2119103. The authors also acknowledge the computational resources provided by Texas A&M High Performance Research Computing (HPRC) Facility.

Conflict of Interest

There are no conflicts of interest.

Data Availability Statement

The datasets generated and supporting the findings of this article are obtainable from the corresponding author upon reasonable request.

References

[1] Rasmussen, C. E., and Williams, C. K. I., 2005, *Gaussian Processes for Machine Learning (Adaptive Computation and Machine Learning)*, The MIT Press, Cambridge, MA.

[2] Shahriari, B., Swersky, K., Wang, Z., Adams, R. P., and De Freitas, N., 2015, "Taking the Human Out of the Loop: A Review of Bayesian Optimization," Proc. IEEE, **104**(1), pp. 148–175.

[3] Frazier, P. I., 2018, "Bayesian Optimization," *Recent Advances in Optimization and Modeling of Contemporary Problems-Informs*, Catonsville, MD, pp. 255–278.

[4] Joy, T. T., Rana, S., Gupta, S., and Venkatesh, S., 2020, "Batch Bayesian Optimization Using Multi-scale Search," Knowl.-Based Syst., **187**(104818).

[5] Couperthwaite, R., Molkeri, A., Khatamsaz, D., Srivastava, A., Allaire, D., and Arròyave, R., 2020, "Materials Design Through Batch Bayesian Optimization With Multisource Information Fusion," JOM, **72**, pp. 1–13.

[6] Ghoreishi, S. F., Friedman, S., and Allaire, D. L., 2019, "Adaptive Dimensionality Reduction for Fast Sequential Optimization With Gaussian Processes," ASME J. Mech. Des., **141**(7), p. 071404.

[7] Constantine, P. G., Dow, E., and Wang, Q., 2014, "Active Subspace Methods in Theory and Practice: Applications to Kriging Surfaces," SIAM J. Sci. Comput., **36**(4), pp. A1500–A1524.

[8] Khatamsaz, D., Molkeri, A., Couperthwaite, R., James, J., Arróyave, R., Srivastava, A., and Allaire, D., 2021, "Adaptive Active Subspace-Based Efficient Multifidelity Materials Design," Mater. Des., **209**.

[9] Zhang, Y., Hoang, T. N., Low, B. K. H., and Kankanhalli, M., 2017, "Information-Based Multi-fidelity Bayesian Optimization," NIPS Workshop on Bayesian Optimization, p. 49.

[10] Ghoreishi, S. F., Molkeri, A., Srivastava, A., Arroyave, R., and Allaire, D., 2018, "Multi-information Source Fusion and Optimization to Realize ICME: Application to Dual-Phase Materials," ASME J. Mech. Des., **140**(11), p. 111409.

[11] Ghoreishi, S. F., and Allaire, D., 2019, "Multi-information Source Constrained Bayesian Optimization," Struct. Multidiscipl. Optim., **59**(3), pp. 977–991.

[12] Khatamsaz, D., Peddareddygari, L., Friedman, S., and Allaire, D. L., 2020, "Efficient Multi-information Source Multiobjective Bayesian Optimization," AIAA Scitech 2020 Forum, Orlando, FL, Jan. 6–10, p. 2127.

[13] Khatamsaz, D., Peddareddygari, L., Friedman, S., and Allaire, D., 2021, "Bayesian Optimization of Multiobjective Functions Using Multiple Information Sources," AIAA J., **59**(6), pp. 1964–1974.

[14] McDannald, A., Frontzek, M., Savici, A. T., Doucet, M., Rodriguez, E. E., Meuse, K., Opsahl-Ong, J., et al. , 2022, "On-the-Fly Autonomous Control of Neutron Diffraction Via Physics-Informed Bayesian Active Learning," Appl. Phys. Rev., **9**(2), p. 021408.

[15] Ziatdinov, M. A., Ghosh, A., and Kalinin, S. V., 2022, "Physics Makes the Difference: Bayesian Optimization and Active Learning Via Augmented Gaussian Process," Mach. Learn.: Sci. Technol., **3**(1), p. 015003.

[16] Chakrabarty, A., Wichern, G., and Laughman, C., 2021, "Attentive Neural Processes and Batch Bayesian Optimization for Scalable Calibration of Physics-Informed Digital Twins". *arXiv preprint arXiv:2106.15502*.

[17] Astudillo, R., and Frazier, P. I., 2021, "Thinking Inside the Box: A Tutorial on Grey-Box Bayesian Optimization," WSC, Phoenix, AZ.

[18] Kandasamy, K., Krishnamurthy, A., Schneider, J., and Póczos, B., 2018, "Parallelised Bayesian Optimisation Via Thompson Sampling," Artificial Intelligence and Statistics, Lanzarote, Spain, PMLR, pp. 133–142.

[19] Folch, J. P., Lee, R. M., Shafei, B., Walz, D., Tsay, C., van der Wilk, M., and Misener, R., 2023, "Combining Multi-fidelity Modelling and Asynchronous Batch Bayesian Optimization," Comput. Chem. Eng., **172**.

[20] González, J., Dai, Z., Hennig, P., and Lawrence, N., 2016, "Batch Bayesian Optimization Via Local Penalization," Artificial Intelligence and Statistics, Cadiz, Spain, PMLR, pp. 648–657.

[21] Ginsbourger, D., Janusevskis, J., and Le Riche, R., 2011, "Dealing with Asynchronicity in Parallel Gaussian Process Based Global Optimization," Ph.D. thesis, Mines Saint-Etienn.

[22] Janusevskis, J., Le Riche, R., Ginsbourger, D., and Girdziusas, R., 2012, "Expected Improvements for the Asynchronous Parallel Global Optimization of Expensive Functions: Potentials and Challenges," Learning and Intelligent Optimization, Paris, France, Springer, pp. 413–418.

[23] Allaire, D., and Willcox, K., 2012, "Fusing Information From Multifidelity Computer Models of Physical Systems," 2012 15th International Conference on Information Fusion, IEEE, pp. 2458–2465.

[24] Ghoreishi, S. F., and Allaire, D. L., 2018, "A Fusion-Based Multi-information Source Optimization Approach Using Knowledge Gradient Policies," 2018 AIAA/ASCE/AHS/ASC Structures, Structural Dynamics, and Materials Conference, Kissimmee, FL, p. 1159.

[25] Thomison, W. D., and Allaire, D. L., 2017, "A Model Reification Approach to Fusing Information From Multifidelity Information Sources," 19th AIAA Non-Deterministic Approaches Conference, p. 1949.
[26] Clyde, M., 2003, "Model Averaging," *Subjective and Objective Bayesian Statistics*, 2nd ed., Wiley-Interscience, Hoboken, NJ.
[27] Clyde, M., and George, E., 2004, "Model Uncertainty," Stat. Sci., **19**(No. 1), pp. 81–94.
[28] Draper, D., 1995, "Assessment and Propagation of Model Uncertainty," J. R. Stat. Soc. Ser. B, **57**(1), pp. 45–97.
[29] Hoeting, J., Madigan, D., Raftery, A., and Volinsky, C., 1999, "Bayesian Model Averaging: A Tutorial," Stat. Sci., **14**(4), pp. 382–417.
[30] Leamer, E., 1978, *Specification Searches: Ad Hoc Inference With Nonexperimental Data*, John Wiley & Sons, New York.
[31] Madigan, D., and Raftery, A., 1994, "Model Selection and Accounting for Model Uncertainty in Graphical Models Using Occam's Window," Am. Stat. Assoc., **89**(428), pp. 1535–1546.
[32] Mosleh, A., and Apostolakis, G., 1986, "The Assessment of Probability Distributions From Expert Opinions With an Application to Seismic Fragility Curves," Risk Anal., **6**(4), pp. 447–461.
[33] Reinert, J., and Apostolakis, G., 2006, "Including Model Uncertainty in Risk-Informed Decision Making," Ann. Nucl. Energy, **33**(4), pp. 354–369.
[34] Riley, M., and Grandhi, R., 2011, "Quantification of Modeling Uncertainty in Aeroelastic Analyses," J. Aircr., **48**(3), pp. 866–873.
[35] Zio, E., and Apostolakis, G., 1996, "Two Methods for the Structured Assessment of Model Uncertainty by Experts in Performance Assessments of Radioactive Waste Repositories," Reliab. Eng. Syst. Saf., **54**(2–3), pp. 225–241.
[36] Julier, S., and Uhlmann, J., 2001, "General Decentralized Data Fusion With Covariance Intersection," *Handbook of Data Fusion*, D. Hall, and J. Llinas, eds., CRC Press, Boca Raton, FL.
[37] Julier, S., and Uhlmann, J., 1997, "A Non-Divergent Estimation Algorithm in the Presence of Unknown Correlations," Proceedings of the 1997 American Control Conference, Albuquerque, NM, pp. 2369–2373.
[38] Geisser, S., 1965, "A Bayes Approach for Combining Correlated Estimates," J. Am. Stat. Assoc., **60**(310), pp. 602–607.
[39] Morris, P., 1977, "Combining Expert Judgments: A Bayesian Approach," Manage. Sci., **23**(7), pp. 679–693.
[40] Winkler, R., 1981, "Combining Probability Distributions From Dependent Information Sources," Manage. Sci., **27**(4), pp. 479–488.
[41] Powell, W. B., and Ryzhov, I. O., 2012, *Optimal Learning*, Vol. 841, John Wiley & Sons, Hoboken, NJ.
[42] Frazier, P., Powell, W., and Dayanik, S., 2009, "The Knowledge-Gradient Policy for Correlated Normal Beliefs," INFORMS J. Comput., **21**(4), pp. 599–613.
[43] Frazier, P. I., Powell, W. B., and Dayanik, S., 2008, "A Knowledge-Gradient Policy for Sequential Information Collection," SIAM J. Control Optim., **47**(5), pp. 2410–2439.
[44] Ghoreishi, S. F., Molkeri, A., Arróyave, R., Allaire, D., and Srivastava, A., 2019, "Efficient Use of Multiple Information Sources in Material Design," Acta Mater., **180**, pp. 260–271.
[45] Khatamsaz, D., Molkeri, A., Couperthwaite, R., James, J., Arróyave, R., Allaire, D., and Srivastava, A., 2021, "Efficiently Exploiting Process-Structure-Property Relationships in Material Design by Multi-information Source Fusion," Acta Mater., **206**(116619).
[46] Molkeri, A., Khatamsaz, D., Couperthwaite, R., James, J., Arróyave, R., Allaire, D., and Srivastava, A., 2022, "On the Importance of Microstructure Information in Materials Design: PSP Vs PP," Acta Mater., **223**(117471).
[47] Bhattacharya, D., 2011, "Metallurgical Perspectives on Advanced Sheet Steels for Automotive Applications," *Advanced Steels*, Springer, Berlin/Heidelberg, pp. 163–175.
[48] Chen, P., Ghassemi-Armaki, H., Kumar, S., Bower, A., Bhat, S., and Sadagopan, S., 2014, "Microscale-Calibrated Modeling of the Deformation Response of Dual-Phase Steels," Acta Mater., **65**, pp. 133–149.
[49] Srivastava, A., Bower, A., Hector, L., Carsley, J., Zhang, L., and Abu-Farha, F., 2016, "A Multiscale Approach to Modeling Formability of Dual-Phase Steels," Modell. Simul. Mater. Sci. Eng., **24**(2), p. 025011.
[50] Voigt, W., 1889, "On the Relation Between the Elasticity Constants of Isotropic Bodies," Ann. Phys. Chem., **274**, pp. 573–587.
[51] Reuß, A., 1929, "Berechnung Der Fließgrenze Von Mischkristallen Auf Grund Der Plastizitätsbedingung Für Einkristalle.," ZAMM-J. Appl. Math. Mech., **9**(1), pp. 49–58.
[52] Bouaziz, O., and Buessler, P., 2002, "Mechanical Behaviour of Multiphase Materials: An Intermediate Mixture Law Without Fitting Parameter," Metall. Res. Technol., **99**(1), pp. 71–77.