



Active Learning and Bayesian Optimization: A Unified Perspective to Learn with a Goal

Francesco Di Fiore¹ · Michela Nardelli¹ · Laura Mainini^{1,2,3}

Received: 10 October 2023 / Accepted: 9 January 2024 / Published online: 23 April 2024
© The Author(s) 2024

Abstract

Science and Engineering applications are typically associated with expensive optimization problem to identify optimal design solutions and states of the system of interest. Bayesian optimization and active learning compute surrogate models through efficient adaptive sampling schemes to assist and accelerate this search task toward a given optimization goal. Both those methodologies are driven by specific infill/learning criteria which quantify the utility with respect to the set goal of evaluating the objective function for unknown combinations of optimization variables. While the two fields have seen an exponential growth in popularity in the past decades, their dualism and synergy have received relatively little attention to date. This paper discusses and formalizes the synergy between Bayesian optimization and active learning as symbiotic adaptive sampling methodologies driven by common principles. In particular, we demonstrate this unified perspective through the formalization of the analogy between the Bayesian infill criteria and active learning criteria as driving principles of both the goal-driven procedures. To support our original perspective, we propose a general classification of adaptive sampling techniques to highlight similarities and differences between the vast families of adaptive sampling, active learning, and Bayesian optimization. Accordingly, the synergy is demonstrated mapping the Bayesian infill criteria with the active learning criteria, and is formalized for searches informed by both a single information source and multiple levels of fidelity. In addition, we provide guidelines to apply those learning criteria investigating the performance of different Bayesian schemes for a variety of benchmark problems to highlight benefits and limitations over mathematical properties that characterize real-world applications.

1 Introduction

In science and engineering, the development of advanced technologies involves the formalization and solution of optimization problems to identify both optimal designs capable to satisfy competing requirements of performance [85], and states of the system to monitor their health status during the operational life [66]. Depending on the specific application, the identification of optimal solutions requires the minimization of an objective function that measures the goodness of design configurations with respect to the requirements, or

the accuracy of the estimated health status of the system as to measurements. Typically, the scale of complexity of engineering systems requires several evaluations of this objective function through accurate computer simulations—e.g. Computational Fluid Dynamics (CFD) or Computational Structural Dynamics (CSD)—or physical experiments—e.g. lab-scale test benches or real-world testing—before assessing an optimal solution. The use of highly complicate representations of those systems leads to a significant bottleneck: the demand for resources to evaluate the objective function for all the combinations of optimization variables is difficult to be adequately satisfied. Indeed, the acquisition of data from these high-fidelity models involves huge non-trivial computational and economical costs that could arise from the computation of the objective function and its derivatives over ideally the entire optimization domain.

Surrogate models are computed on evaluations of the objective function acquired through computer codes and/or physical experiments of the system: these sources of information are mostly treated as purely input/output black-box

✉ Francesco Di Fiore
francesco.difiore@polito.it

¹ Department of Mechanical and Aerospace Engineering, Politecnico di Torino, Turin, Italy

² Department of Aeronautics, Imperial College London, London, United Kingdom

³ Brahma Vasudevan Institute for Sustainable Aviation, London, United Kingdom

relationship whose analytical form is unknown and not directly accessible to the optimizer. Thus, the accuracy and efficiency of the resulting surrogate are highly dependent on the sampling approach adopted to select informative combinations of optimization variables for the acquisition of data. Among the numerous sampling schemes available in literature, it is possible to identify two major families: one-shot, and sequential schemes. The one-shot strategy defines a grid of samples over the domain all at once. Examples include Latin Hypercube [88], factorial and fractional factorials designs [42, 94], Plackett-Burmann [44], and D-optimal [95]. However, it is very hard to identify a priori the best design of those experiments to efficiently compute the most informative surrogate. To overcome these limitations, sequential sampling selects samples over the domain through an iterative process [16, 59]. Among these, adaptive sampling [108] provides resource-efficient techniques that seek to reduce as much as possible the evaluations of the objective function, and target the improvement of the fitting quality across the domain and/or the acceleration of the optimization search [24, 79, 133]. Popular adaptive samplings to address black-box optimization problems characterized by the expensive evaluation of the objective function are those realized through the Bayesian Optimization (BO) methodology [33, 124]. BO aims at efficiently elicit valuable data from models of the system to contain the computational expense of the optimization procedure. The Bayesian routine iteratively computes a surrogate model of the objective function, and defines a goal-driven sampling process through an acquisition function computed on the surrogate information. This acquisition function measures the merit of samples according to certain infill criteria, and permits to select the next sample that maximizes the query utility with respect to the given optimization goal.

The popular paradigms for Bayesian optimization show substantial synergy with active learning schemes which has not been explicitly discussed and formally described in literature to date. This paper proposes the explicit formalization of this synergy through an original perspective of Bayesian optimization and active learning as symbiotic expressions of adaptive sampling schemes. The aim of this unifying viewpoint is to support the use of those methodologies, and point out and discuss the analogies via their mathematical formalization. This unified interpretation is based on the formulation and demonstration of the analogy between the Bayesian infill criteria and the active learning criteria as the elements responsible for the decision on how learn from samples to reach the given goal. In support of this unified perspective, this paper first clarifies the concept of goal-driven learning, and proposes a general classification of adaptive sampling methods that recognizes Bayesian optimization and active learning as methodologies characterized by goal-oriented search schemes. Thus, we elucidate the

synergy between Bayesian optimization and active learning mapping the Bayesian learning features on the active learning properties. The mapping is discussed through the analysis of three popular Bayesian frameworks for both the case of a single information source, and when a spectrum of multiple sources are available to the search. In addition, we observe the capabilities introduced by the different learning criteria over a comprehensive set of benchmark problems specifically defined to stress test and validate goal-driven approaches [83]. The objective is to discuss opportunities and limitations of different learning principles over a variety of challenging mathematical properties of optimization problems frequently encountered in complex scientific and engineering applications.

This manuscript is organized as follows. Section 2 discusses goal-driven learning procedures and defines the concept of goal-driven learner according to surrogate modeling and optimization. In Sect. 3, we recognize that Bayesian optimization, active learning and adaptive sampling are not fully superimposable concepts, and propose a general classification to position Bayesian optimization and active learning with respect to the adaptive sampling methodologies. Then, Sect. 4 provides an overview on Bayesian optimization and multifidelity Bayesian optimization. Section 5 presents our perspective on the symbiotic relationship between Bayesian optimization and active learning. Then, in Sect. 6 popular Bayesian optimization and multifidelity Bayesian optimization algorithms are numerically investigated over a variety of benchmark problems. Finally, Sect. 7 provides concluding remarks.

2 Goal-Driven Learning

Goal-driven learning is a decision-making process in which each decision is made to acquire specific information about the system of interest that contributes the most to achieve a given goal [11, 21, 40, 78, 99, 109]. This learning goal can be the increase of the knowledge of the system behaviour over all the domain of application, or acquire specific knowledge to enhance and accelerate the identification of optimization solutions. Accordingly, a goal-driven learner selects what to learn considering both the current knowledge and information needed, and determines how to learn quantifying the relative utility of alternative options in the current circumstances.

This paper focuses on Bayesian optimization and active learning as goal-driven procedures where a surrogate model is built to accurately represent the behaviour of a system or effectively inform an optimization procedure to minimize given objectives. This goal-driven process is guided by learning principles that determine the “best” location of the

domain to acquire information about the system, and refine the surrogate model towards the goal—improve the accuracy of the surrogate or minimize an objective function over the domain. Formally, these surrogate based modeling and optimization problems can be formulated as a minimization problem of the following form:

$$x^* = \arg \min_{x \in \chi} f(\mathbf{R}(x)) \quad (1)$$

where $f(\mathbf{R}(x))$ denotes the objective function evaluated at the location $x \in \chi$ of the domain χ . The objective function is of the general form $f = f(\mathbf{R}(x))$, where $\mathbf{R}(x)$ represents the response of the system of interest evaluated through a model—e.g. computer-based numerical simulations or real-world experiments. In surrogate based modeling, the objective function can be represented as the error between the approximation of the surrogate model and the response of the system: the goal is to minimize such error to improve the accuracy of the surrogate over all the domain. In surrogate based optimization, the objective function represents a performance indicator dependent on the system response: the goal is to minimize this indicator to improve the capabilities of the system according to given performance requirements. Goal-driven techniques address Equation (1) through a decision-making iterative process where learning principles tailor the acquisition of specific knowledge about the objective function—evaluation of f at certain domain location x —currently needed to update the surrogate and inform the learner towards the given goal.

In this context, the goal-driven learner is the agent that makes decisions based on the current knowledge of the system of interest, and acquires new information to accomplish a given goal while augmenting the awareness about the system itself. In practice, the learner queries the sample that maximizes the utility to achieve the desired goal: specific learning principles quantify this utility based on the surrogate estimate and in response to information needs. At the same time, the surrogate model is dynamically updated

once new information are acquired, and informs the learner to focus and tailor on the fly the elicitation of samples to further overarching the goal. Thus, the distinguishing element of a goal-driven learning procedure is represented by the mutual exchange of information between the learner and the surrogate model: the learner assimilates the information from the surrogate to make a decision aimed at achieving the goal, and the approximation/prediction of the surrogate is enriched by the result of this decision.

3 Adaptive Sampling Classification

Bayesian optimization and active learning realize adaptive sampling schemes to efficiently accomplish a given goal while adapting to the previously collected information. In recent years, there has been a profusion of literature devoted to the general topic of adaptive sampling but arguably a blurring of focus: many contributions from different field provided a deal of interesting advancements, but also led to some degree of confusion around the concepts of adaptive sampling, active learning and Bayesian optimization. Figure 1 illustrates the use of the words “adaptive sampling”, “active learning”, and “Bayesian optimization” from 1990 to 2022. In addition, we report the combined use of all the three words over the same period of time. It can be appreciated both the general increasing trend of use of the three techniques and the associated increase of the use of the three terms combined. Many times the three concepts have been used as complete synonyms, with some growing abuse motivated by the difficulties to map the (shaded) boundaries.

Stemming from these considerations, this paper recognizes that adaptive sampling is not always superimposable with active learning and Bayesian optimization. Figure 2 illustrates the relationships between those three methodologies. We propose a classification of adaptive sampling techniques in three main families, namely adaptive probing (Sect. 3.1), adaptive modeling (Sect. 3.2) and adaptive learning (Sect. 3.3). This classification is based on the concept of goal-driven learning as the distinctive element of adaptive learning methodologies: the learner assimilates the information from the surrogate model to make a decision aimed at achieving a goal, and the surrogate is enriched by the result of this decision following a mutual exchange of information. Conversely, adaptive probing and adaptive modeling classes do not realize a goal-driven learning: the former does not rely on a surrogate model to assist the sampling procedure while the latter computes a surrogate model that is not used to inform the search task. This classification permits to clarify the reciprocal positions between adaptive sampling, active learning and Bayesian optimization.




Fig. 1 Citations of Bayesian Optimization (BO), Active Learning (AL), Adaptive Sampling (AS) and the three terms combined (BO+AL+AS)




Fig. 2 Where adaptive sampling and active learning meet: this work focuses on the synergies between Bayesian optimization and active learning as goal-driven learning procedures driven by common learning principles

Accordingly, adaptive sampling and active learning do not completely overlap. Active learning strategies are categorized into population-based and pool-based algorithms according to the nature of the search procedure [129, 141]. In population-based active learning, the distribution of the objective function is available: the learner seeks to determine the optimal training input density to generate training points without relying on a surrogate model of the objective function. Conversely, pool-based active learning computes a surrogate model of the unknown objective function that is used to inform the learner toward a given goal, and is updated during the procedure to refine the informative content supporting the learning procedure. Thus, pool-based active learning methods realize goal-driven learning schemes and can be collocated in the adaptive learning class while population-based active learning techniques can not be considered as adaptive samplings. Following the proposed classification, Bayesian optimization represents the logic intersection between active learning and adaptive sampling since (i) BO realizes an adaptive sampling scheme towards a given goal, and (ii) the BO goal-driven learning procedure is guided by learning principles also traceable in active learning schemes. This synergy between Bayesian optimization and active learning is the main focus of our work, and the remaining of this manuscript is dedicated to formalize and discuss this dualism. To support this discussion, we provide additional details of the proposed classification for adaptive sampling, and review some popular approaches for each of the three classes. The literature on adaptive sampling is vast, and a complete review goes beyond the purpose of this work. Although our discussion will not be comprehensive, the objective is to highlight the distinguishing features of each class and clarify the relative positions of adaptive sampling, active learning and Bayesian optimization.

3.1 Adaptive Probing

Adaptive probing schemes exploit the observations of previous samples without computing any surrogate model. These sampling procedures are informed exclusively from the collected data to guide the selection of the next location to query, and exclude the adoption of emulators to support the search. Several adaptive probing frameworks have been developed based on the Monte Carlo method [101, 125]. Among these, adaptive importance samplings [10, 64, 102] and adaptive Markov Chain Monte Carlo samplings [2, 3] represent popular methodologies adopted in different practical scenarios, from signal processing [9, 159] to reliability analysis of complex systems [58, 147]. Adaptive importance sampling uses previously observed samples to adapt the proposal densities and locate the regions from which samples should be drawn; this strategy permits to iteratively improve the quality of the samples distribution and enhance the accuracy of the relative inference from these observations. Adaptive Markov Chain Monte Carlo (MCMC) determines the parameters of the MCMC transition probabilities on the fly through already collected information. This adaptively generates new samples from an usually complex and high-dimensional distribution, and enhances the overall computational efficiency and reliability of the procedure. In the next paragraph, we report the mathematical formulation of adaptive importance sampling to illustrate the properties of adaptive probing methodologies and the elements that differentiate them from active learning paradigms.

3.1.1 Adaptive Importance Sampling

Adaptive Importance Sampling (AIS) usually considers a generic inference problem characterized by a certain probability density function (pdf) $\tilde{\pi}(\mathbf{x})$ of a d_x -dimensional vector of unknown statistic real parameters $\mathbf{x} \in \chi$. AIS frameworks aim to provide a numerical approximation of some particular moment of \mathbf{x} :

$$I(f) = \mathbb{E}_{\tilde{\pi}}[f(\mathbf{x})] = \int f(\mathbf{x})\tilde{\pi}(\mathbf{x}) d\mathbf{x} \quad (2)$$

where $f : \chi \rightarrow \mathbb{R}$ can be any function of \mathbf{x} integrable with respect to the pdf $\tilde{\pi}(\mathbf{x})$

The integral $I(f)$ is representative of different mathematical problems, from Bayesian inference [113] to the estimate of rare events [48]. In many practical scenarios, the integral $I(f)$ cannot be computed in closed form. Adaptive importance sampling provides an algorithmic framework to efficiently address this problem.

Let us define a proposal probability density function $q(\mathbf{x})$ to simulate samples under the restriction that $q(\mathbf{x}) > 0$

for all \mathbf{x} where $\tilde{\pi}(\mathbf{x})f(\mathbf{x}) \neq 0$. AIS provides an iterative procedure that improves the quality of one or multiple proposals $q(\mathbf{x})$ to approximate a non-normalized non-negative target function $\pi(\mathbf{x})$. At the beginning, AIS initializes N proposals $\{q_n(\mathbf{x}|\theta_{n,1})\}_{n=1}^N$ parameterized through the vector $\theta_{n,1}$. Then, the procedure simulates K samples from each proposal $\mathbf{x}_{n,1}^{(k)}$, $n = 1, \dots, N$, $k = 1, \dots, K$, and assigns to each sample an associated importance weight formalized as follows:

$$w_n = \frac{\pi(\mathbf{x}_n)}{q(\mathbf{x}_n)}, \quad n = 1, \dots, N \quad (3)$$

These importance weights measure the representativeness of each sample simulated from the proposal pdf $q(\mathbf{x})$ with reference to the distribution of random variables $\tilde{\pi}(\mathbf{x})$.

At this point, this set of N weighted samples $\{\mathbf{x}_{n,1}^{(k)}, w_{n,1}^{(k)}\}$, $n = 1, \dots, N$, $k = 1, \dots, K$ are used to define a self-normalized estimator:

$$\hat{I}^N(f) = \sum_{n=1}^N \bar{w}_n f(w_n) \quad (4)$$

where $\bar{w}_n = w_n / \sum_{j=1}^N w_j$ are the normalized weights. This permits to approximate the target function distribution as follows:

$$\tilde{\pi}^N(\mathbf{x}) = \sum_{n=1}^N \bar{w}_n \delta(\mathbf{x} - \mathbf{x}_n) \quad (5)$$

where δ represents the Dirac measure.

Finally, AIS realizes the adaptation phase and updates the parameters of the n -th proposals from $\theta_{n,1}$ to $\theta_{n,2}$ using the last set of drawn parameters [84] or all the parameters evaluated so far [29]. The whole procedure is repeated until a certain termination criteria is met (e.g. maximum number of iterations).

This adaptive policy permits to gradually evolve the single or multiple proposal densities to accurately approximate the target pdf. The generation of new samples is uniquely driven by the measurement of the importance of previous samples (weighting) that supports the updating of the proposal parameters (adaptation). Thus, AIS adaptively locates promising regions to query without benefit from an overall quantification of the goodness of all the spectrum of samples available in the domain—e.g. through the construction of a surrogate model. On this basis, AIS and the general class of adaptive probing strategies are not considerable as learning procedures since the adaptation phase is not informed by a surrogate model updated on the fly during the procedure, and is not guided by a “learner” that assimilates information from this emulator and adapts the next queries to achieve a given goal.

3.2 Adaptive Modeling

Adaptive modeling paradigms sample the domain supported by the information from previous queries, and use the collected data to build a surrogate model. However, the informative content encoded in the emulator is not used to guide the sampling and decide the next point to evaluate. Adaptive modeling approaches have been extensively developed for the reliable propagation and quantification of uncertainties [56, 57], analysis of ordinary or partial differential equations [27, 43], and inverse problems [82, 86]. One common approach is represented by adaptive stochastic collocation methodologies, which uses an adaptive sparse grid approximation scheme to construct an interpolant polynomial in a multi-dimensional random space [45, 72]. The adaptive selection of collocation points is driven by an error indicator [37] or estimator [41] that evaluates a certain number of sparse admissible subspaces of the domain: the subspace that exhibits the higher error is included in the grid and the new set of subspaces is identified. Other well-known adaptive modeling approaches are residual-based samplings distribution [146]. This family of techniques is mostly applied to improve the training efficiency of Physics-Informed Neural Networks (PINN) surrogate models. Residual-based approaches enhance the distribution of residual points by placing more samples according to certain properties of the residuals during the training of PINN. This decision can be made on the basis of locations where the residual of the partial differential equation is large [81], according to a probability density function of the residual points [96], and hybrid approaches of the above [146]. This permits to achieve better accuracy of the final PINN surrogate model while containing the computational burden associated with computations. Both stochastic collocation and residual based samplings are intended to build an efficient and accurate surrogate model over the domain of samples. However, the sampling procedure is adapted uniquely to previous evaluated samples without a learning procedure from data: the surrogate model is not used to inform the decision on where to sample, and is not progressively updated with previous information. In the following, we provide general mathematical details about adaptive stochastic collocation to analyze the peculiarities of the adaptive modeling class, and underline the absence of a learning process during the construction of the surrogate model.

3.2.1 Adaptive Stochastic Collocation

Adaptive Stochastic Collocation (ASC) builds an interpolation function to approximate the outputs from a model of interest. This emulator is constructed on the evaluations of

the model at valuable collocation points of the stochastic inputs to obtain the moments and the probability density function of the outputs.

Consider any point \mathbf{x} contained in the random space $\Gamma \subset \mathbb{R}^N$ with probability distribution function $\rho(\mathbf{x})$. The goal of ASC is to find an interpolating polynomial $\mathcal{I}(f)$ to approximate a smooth function $f(\mathbf{x}) : \mathbb{R}^N \rightarrow \mathbb{R}$:

$$\mathcal{I}(f)(\mathbf{x}_k) = f(\mathbf{x}_k), \quad 1 \leq k \leq P \tag{6}$$

for a given set of points $\{\mathbf{x}_k\}_{k=1}^P$. The selection of the collocation points majorly influences the capability of the interpolating polynomial to be close to the original function f . For multivariate problems, the interpolation function is defined as follows using the tensor product grid:

$$\begin{aligned} \mathcal{I}(f) &= (\mathcal{U}^{i_1} \otimes \dots \otimes \mathcal{U}^{i_N})(f) \\ &= \sum_{j_1=1}^{n_{i_1}} \dots \sum_{j_N=1}^{n_{i_N}} f(\mathbf{x}_{j_1}^{i_1}, \dots, \mathbf{x}_{j_N}^{i_N}) \cdot (\mathcal{L}_{j_1}^{i_1} \otimes \dots \otimes \mathcal{L}_{j_N}^{i_N}) \end{aligned} \tag{7}$$

where \mathcal{U}^{i_k} is the univariate interpolation function for the level i_k in the k -th coordinate, $\mathbf{x}_{j_m}^{i_m}$ is the j_m -th node, and \mathcal{L}_{j_k} are the Lagrange interpolating polynomials.

Equation 7 demands for $n_{i_1} \times \dots \times n_{i_N}$ nodes, which indicates an exponential rate of computational cost growth with the number of dimensions. Adaptive stochastic collocation targets the reduction of this computational effort through an adaptive sparse grid of collocation points: the objective is to wisely place more points of the grid in the important directions to prioritize the collection of highly informative data. This adaptive sparse grid is defined through a subset of the full tensor product grid as follows:

$$\begin{aligned} \mathcal{A}_{q,N}(f) &= \sum_{|\mathbf{i}| \leq q} (\Delta \mathcal{U}^{i_1} \otimes \dots \otimes \Delta \mathcal{U}^{i_N})(f) \\ &= \mathcal{A}_{q-1,N}(f) + \sum_{|\mathbf{i}|=q} (\Delta \mathcal{U}^{i_1} \otimes \dots \otimes \Delta \mathcal{U}^{i_N})(f) \end{aligned} \tag{8}$$

where $\mathbf{i} = (i_1, \dots, i_N) \in \mathbb{R}^N$, $|\mathbf{i}| = i_1 + \dots + i_N$, q is the sparseness parameter, and the difference formulas are defined by $\mathcal{U}^0 = 0$ and $\Delta \mathcal{U}^i = \mathcal{U}^i - \mathcal{U}^{i-1}$.

Equation 8 leverages the previous results to extend the interpolation from level $q - 1$ to q through the evaluation of the multivariate function on the sparse grid:

$$\begin{aligned} \mathcal{H}_{q,N} &= \bigcup_{|\mathbf{i}| \leq q} (\Delta \vartheta^{i_1} \times \dots \times \Delta \vartheta^{i_N}) \\ &= \mathcal{H}_{q-1,N} + \bigcup_{|\mathbf{i}|=q} (\Delta \vartheta^{i_1} \times \dots \times \Delta \vartheta^{i_N}) \end{aligned} \tag{9}$$

where $\Delta \vartheta^i = \vartheta^i \setminus \vartheta^{i-1}$ are the newly added set of univariate nodes ϑ^{i_k} for level i_k in the k -th coordinate.

This scheme adapts the sampling procedure through the knowledge acquired on the fly, and efficiently leverages data

to improve the quality of the interpolation function. In this case, the selection of the collocation points is intended to compute an emulator of the target function, but the adaptive sampling is not driven by the information acquired from this emulator. In addition, the acquisition of data is not used to learn and update the surrogate model. These considerations on ASC can be extended to the general class of adaptive modeling methods: even if the sampling scheme is conceived to construct surrogate models, the selection of promising locations to query is not delegated to a goal-driven learner that leverages a mutual exchange of information with the surrogate.

3.3 Adaptive Learning

Adaptive learning methodologies realize goal-driven learning processes characterized by the mutual exchange of information between the surrogate model and the goal-driven learner: the former is updated and refined after new evaluations of samples while the latter decides the next query based on the updated approximation given by the emulator. Bayesian optimization and pool-based active learning belong to this specific class of adaptive sampling techniques. Bayesian frameworks constitute a learning process driven by the mutual informative assimilation between an acquisition function—learner—and a surrogate model [33, 92]. The acquisition function commensurates the benefit of evaluating samples based on the prediction of the surrogate model, and selects the most useful sample to query towards the given goal—either improve the accuracy of the surrogate over the domain or effectively inform the optimization search; at the same time, the emulator is enriched with the data from the new query, and is updated to refine the approximation of the objective function over the domain. Similarly, pool-based active learning methods search the domain through a goal-driven learner informed by a classification model of samples [120, 153]. This process is characterized by the reciprocal flow of information between the learner and the emulator: the classification model is updated through the new evaluations of unsampled locations, and the learner uses this information to select the next query. Mathematical details about pool-based active learning are provided in the following section to better clarify the distinction between this class of adaptive learners, and the other classes which do not realize a goal-driven learning procedure.

3.3.1 Pool-Based Active Learning

Pool-based active learning commonly defines an optimal sampling strategy to improve the accuracy of a surrogate model adopted to classify data-points from a target distribution of labels over the domain of samples χ . Considering this general classification task, pool-based active learning

routine is grounded on a probabilistic estimate of the distribution of features f over the entire domain χ through a surrogate model \hat{f} . This emulator is trained on a set of collected data-points, and maps features to labels $f_N(\mathbf{x}_n) = \hat{f}_n$ through a predicted probability $p_N(f_n = f|\mathbf{x}_n)$ that estimates the distribution of features over the domain. Suppose we have collected from a large pool of unlabelled data χ the—small— dataset $\mathcal{D}_N\{\mathbf{x}_n, f(\mathbf{x}_n)\}_{n=1}^N$ observing the label values $f(\mathbf{x}_n)$ in output from an observation model or oracle at some informative locations \mathbf{x}_n . Based on this dataset, the goal-driven procedure learns a surrogate model \hat{f}_N whose predictive framework emulates the behaviour of samples over the domain based on the previous collected information.

At this point, an utility function acts as the goal-driven learner informed by the surrogate model, and identifies the most promising sample to be labelled by the oracle according to a measure of utility with respect to the given goal—improve the accuracy of the classifier. The next query augments the dataset $\mathcal{D}_{N+1} = \mathcal{D}_N \cup \{\mathbf{x}_{N+1}, f_{N+1}\}$ and the surrogate model is updated. This utility function defines a learning policy that maps the current predictive distribution to a decision/action on where to sample in the next iteration as follows:

$$\mathbf{x}_{N+1} = \arg \max U(p_N(f_n = f|\mathbf{x}_n)) \quad (10)$$

Equation (10) mathematically formalizes the concept of goal-driven learning procedure: the learner leverages the predicted probability of the surrogate $p_N(y_n = y|\mathbf{x}_n)$ to make an action \mathbf{x}_{N+1} ; at the same time, the decision is used to enrich the dataset $\mathcal{D}\{\mathbf{x}_n, f(\mathbf{x}_n)\}_{n=1}^{N+1}$ and update the predicted probability p_{N+1} . This mutual exchange and assimilation between the learner and the surrogate represents the key aspect that defines a goal-driven learning process and the whole class of adaptive learning sampling schemes.

4 Bayesian Frameworks

Bayesian optimization constitutes the mid-point between adaptive sampling and active learning. This intersection represents the focal point of our work, and motivates the substantial synergy between Bayesian optimization and active learning as adaptive sampling schemes capable to learn from data and accomplish a certain learning goal. The remaining of this section is dedicated to the general overview of Bayesian optimization considering both a single source of information (Sect. 4.1) and when multiple sources are available to the learning procedure (Sect. 4.2). This will guide the reader into the next sections that make explicit the symbiosis between Bayesian frameworks and active learning through our original perspective of Bayesian optimization as a way to actively learn with acquisition functions (Sect. 5).

4.1 Bayesian Optimization

The birth of Bayesian optimization can be retraced in 1964 with the work of Kushner [69] where unconstrained one-dimensional optimization problems are addressed through a predictive framework based on the Wiener process surrogate model, and a sampling scheme guided by the probability of improvement acquisition function. Further contributions have been proposed by Zhilinskas [158] and Mockus [90], and the methodology has been extended to high dimensional optimization problems in the works of Stuckman [128] and Elder [28]. Bayesian optimization achieved resounding success after the introduction of the Efficient Global Optimization (EGO) algorithm by Jones et al. [61]. EGO uses a Kriging surrogate model to predict the distribution of the objective function, and adopts the expected improvement acquisition function to measure the improvement of the optimization procedure obtained evaluating unknown samples.

The EGO methodology paves the way to the application of Bayesian optimization over a wide range of problems in science and engineering. These research fields demand for the efficient management of the information from black-box representations of the objective function—the procedure is only aware of the input and output without a priori knowledge about the function—to guide the optimization search. Engineering has been pioneering in the adoption of Bayesian optimization: the design optimization of complex systems is frequently characterized by computationally intensive black-box functions which require efficient global optimization methods. Early applications relates to engineering design optimization [152], computer vision [143] and combinatorial problems [75]. Nowadays, the Bayesian framework becomes widely adopted in many fields including and not limited to engineering [34, 67, 71, 107], robotics and reinforcement learning [4, 7, 150], finance and economics [39, 106], automatic machine learning [132, 134], and preference learning [26, 68]. In addition, significant advances have been made in the expansion of BO methodologies to higher-dimensional search spaces frequently encountered in science and engineering, where the effectiveness of the search procedure is usually correlated to an exponential growth of the required observations of the objective function and associated demand for computational resources and time. Within this context, BO techniques have been scaled to approach high-dimensional problems exploiting potential additive structures of the objective function [62, 138], mapping high-dimensional search spaces into low-dimensional subspaces [97, 137], learning from observations of multiple input points evaluated through parallel computing [123, 139], and through simultaneous local optimization approaches [30].

Given a black-box expensive objective function $f : \chi \rightarrow \mathbb{R}$, Bayesian optimization seeks to identify the input $x^* \in \min_{x \in \chi} f(x)$ that minimizes the objective f over

an admissible set of queries χ with a reduced computational cost. To achieve this goal, Bayesian optimization relies on an adaptive learning scheme based on a surrogate model that provides a probabilistic representation of the objective f , and uses this information to compute an acquisition function $U(\mathbf{x}) : \chi \rightarrow \mathbb{R}^+$ that drives the selection of the most promising sample to query. Let us consider the available information regarding the objective function f stored in the dataset $\mathcal{D}_N = \{(\mathbf{x}_1, y_1), \dots, (\mathbf{x}_n, y_n)\}$ where $y_n \sim \mathcal{N}(f(\mathbf{x}_n), \sigma_\epsilon(\mathbf{x}_n))$ are the noisy observations of the objective function and σ_ϵ is the standard deviation of the normally distributed noise.

At each iteration of the optimization procedure, the surrogate model depicts possible explanations of f as $f \sim p(f|\mathcal{D}_N)$ applying a joint distribution over its behaviour at each sample $\mathbf{x} \in \chi$. Typically, Gaussian Processes (GPs) have been widely used as the surrogate model for Bayesian optimization [100, 110]. In GP regression, the prior distribution of the objective $p(f)$ is combined with the likelihood function $p(\mathcal{D}_N|f)$ to compute the posterior distribution $p(f|\mathcal{D}_N) \propto p(\mathcal{D}_N|f)p(f)$, representing the updated beliefs about f . The GP posterior is a joint Gaussian distribution $p(f|\mathcal{D}_N) = \mathcal{N}(\mu(\mathbf{x}), \kappa(\mathbf{x}, \mathbf{x}'))$ completely specified by its mean $\mu(\mathbf{x}) = \mathbb{E}[f(\mathbf{x})]$ and covariance (also referred as kernel) function $\kappa(\mathbf{x}, \mathbf{x}') = \mathbb{E}[(f(\mathbf{x}) - \mu(\mathbf{x}))(f(\mathbf{x}') - \mu(\mathbf{x}'))]$, where $\mu(\mathbf{x})$ represents the prediction of the GP model at \mathbf{x} and $\kappa(\mathbf{x}, \mathbf{x}')$ the associated uncertainty.

BO uses this statistical belief to make the decision of where to sample assisted by an acquisition function U , which identifies the most informative sample $\mathbf{x}_{new} \in \chi$ that should be evaluated via maximization $\mathbf{x}_{new} \in \max_{\mathbf{x} \in \chi} U(\mathbf{x})$. Then, the objective function is evaluated at \mathbf{x}_{new} and this information is used to update the dataset $\mathcal{D}_N = \mathcal{D}_N \cup (\mathbf{x}_{new}, y(\mathbf{x}_{new}))$. Acquisition functions are designed to guide the search for the optimum solution according to different infill criteria which provide a measure of the improvement that the next query is likely to provide with respect to the current posterior distribution of the objective function. In engineering applications, we could retrieve different implementations proposed for the acquisition function, which differ for the infill schemes adopted to sample pursuing the optimization goal. Examples include the Probability of Improvement (PI) [69], Expected Improvement (EI) [61], Entropy Search (ES) [47] and Max-Value Entropy Search (MES) [135], Knowledge-Gradient (KG) [116], and non-myopic acquisition functions [70, 142].

The Probability of Improvement (PI) acquisition function encourages the selection of samples that are likely to obtain larger improvements over the current minimum predicted by the surrogate model, while the Expected Improvement (EI) considers not only the PI but also the expected gain in the solution of the optimization problem achieved evaluating a certain sample. Other popular schemes are entropy-based acquisition functions such as the Entropy Search (ES) and

Max-Value Entropy Search (MES), which rely on estimating the entropy of the location of the optimum and the minimum function value, respectively, to maximize the mutual information between the samples and the location of the global optimum. Knowledge-gradient sampling procedures are conceived for applications where the evaluations of the objective function are affected by noise, recommending the location that maximizes the increment of the expected value that would be acquired by taking a sample from the location. Through the adoption of non-myopic acquisition functions, the learner maximizes the predicted improvement at future iterations of the optimization procedure, overcoming myopic schemes where the improvement of the solution is measured at the immediate step ahead.

4.2 Multifidelity Bayesian Optimization

The evaluation of black-box functions in engineering and science frequently requires time-consuming lab experiments or expensive computer-based models, which would dramatically increase the computational burden for the optimization procedure. This is the case of large-scale design optimization problems, where the evaluation of the objective function for enough samples can not be afforded in practice. In many real-world applications, the objective function can be computed using multiple representations at different levels of fidelity $\{f^{(1)}, \dots, f^{(L)}\}$, where the lower the level of fidelity the less accurate but also less time-consuming the evaluation procedure. Multifidelity methods recognize that different representative levels of fidelity and associated cost can be used to accelerate the optimization process, and enable a flexible trade-off between computational cost and accuracy of the solution. In particular, multifidelity optimization leverages low-fidelity data to massively query the domain, and uses a reduced number of high-fidelity observations to refine the belief about the objective function toward the optimum [6, 32, 103].

Accordingly, Multifidelity Bayesian Optimization (MFBO) learns a surrogate model that synthesizes through stochastic approximation the multiple levels of fidelity available, and uses an acquisition function as the learner that selects the most promising sample and associated level of fidelity to interrogate. This learning procedure provides potential accelerations of the optimization procedure that is reflected in the likely improvement of the surrogate accuracy. According to Godino et al. [38], the improvement in performance occurs usually if the acquisition of large amount of high-fidelity is hampered by the computational expense, the correlation between high-fidelity and low-fidelity data is high, low-fidelity models are sufficiently inexpensive; Under different circumstances, multifidelity optimization might not deliver substantial accelerations and quality of the surrogate: the relationship between dimension of the training set

and surrogate accuracy is not monotonically increasing, as evidenced by [19]. In recent years, multifidelity Bayesian optimization has been successfully adopted for optimization problems ranging from engineering design optimization [8, 22, 23, 40, 89, 119], automatic machine learning [63, 145], applied physics [55, 140], and medical applications [104, 105]. In the context of high-dimensional problems, multifidelity Bayesian optimization capitalizes from fast low-fidelity models to alleviate the computational burden associated with the required numerous observations of the objective function to effectively direct the search toward the given goal, and achieved promising results in terms of accuracy and efficiency for applications in quantum control [73], aerospace engineering [115], and reinforcement learning [53].

Multifidelity Bayesian optimization determines a learning procedure informed by the surrogate model of the objective function constructed on the dataset of noisy objective observations $\mathcal{D}_N = \{(\mathbf{x}_1, y_1^{(l_1)}), \dots, (\mathbf{x}_n, y_n^{(l_n)})\}$, where $y_n^{(l_n)} \sim \mathcal{N}(f^{(l_n)}(\mathbf{x}_n), \sigma_\epsilon(\mathbf{x}_n))$ and σ_ϵ have the same distribution over the fidelities. This multifidelity surrogate model defines an approximation of the objective $f^{(l)} \sim p(f^{(l)}|\mathbf{x}, l, \mathcal{D}_N)$ at different level of fidelity, and represents the belief about the distribution of the objective function over the domain \mathcal{X} based on data. A popular practice for MFBO is to extend the Gaussian process surrogate model to a multifidelity setting through an autoregressive scheme [65]:

$$f^{(l)} = \rho f^{(l-1)}(\mathbf{x}) + \zeta^{(l)}(\mathbf{x}) \quad l = 2, \dots, L \quad (11)$$

where ρ is a constant scaling factor that includes the contribution of the previous fidelity with respect to the following one, and $\zeta^{(l)} \sim GP(0, \kappa^{(l)}(\mathbf{x}, \mathbf{x}'))$ models the discrepancy between two adjoining levels of fidelity. The posterior of the multifidelity Gaussian process is completely specified by the multifidelity mean function $\mu^{(l)}(\mathbf{x}, l) = \mathbb{E}[f^{(l)}(\mathbf{x})]$ that represents the approximation of the objective function at different levels of fidelity, and the multifidelity covariance function $\kappa^{(l)}((\mathbf{x}, l), (\mathbf{x}', l)) = \mathbb{E}[(f^{(l)}(\mathbf{x}, l) - \mu^{(l)}(\mathbf{x}, l))(f^{(l)}(\mathbf{x}', l) - \mu^{(l)}(\mathbf{x}', l))]$ that defines the associated uncertainty for each level of fidelity.

The availability of multiple representations of the objective function poses a further decision task that has to be accounted by the learner during the sampling of unknown locations: the selection of the most promising sample is effected with the simultaneous designation of the information source to be evaluated. This is obtained through a learner represented by the multifidelity acquisition function $U(\mathbf{x}, l)$ that extends the infill criteria of Bayesian optimization, and selects the pair of sample and the associated level of fidelity to query $(\mathbf{x}_{new}, l_{new}) \in \max_{\mathbf{x} \in \mathcal{X}, l \in \mathcal{L}} U(\mathbf{x}, l)$ that is likely to provide higher gains with a regard for the computational expenditure. Among different formulations, well

known multifidelity acquisition functions to address optimization problems are the Multifidelity Probability of Improvement (MFPI) [114], Multifidelity Expected Improvement (MFEI) [51], Multifidelity Predictive Entropy Search (MFPEs) [154], Multifidelity Max-Value Entropy Search (MFMEs) [130], and non-myopic multifidelity expected improvement [21]. These formulations of the acquisition function define adaptive learning schemes that retain the infill principles characterizing the single-fidelity counterpart, and account for the dual decision task balancing the gains achieved through accurate queries with the associated cost during the optimization procedure.

5 An Active Learning Perspective

Bayesian frameworks and Active learning schemes exhibit a strong synergy: in both cases the learner seeks to design an efficient sampling policy to accomplish the learning goal, and is guided by a surrogate model that informs the learner and is continuously updated during the learning procedure. Active learning literature is vast and include a multitude of approaches [1, 12, 14, 20, 111, 121, 122, 141]. According to the well accepted classification proposed by Sugiyama and Nakajima [129], active learning strategies can be categorized in population-based and pool-based active learning frameworks according to the nature of the sampling scheme defined by the learner. Population-based active learning targets the identification of the best optimal density of the samples for training known the target distribution. Conversely, pool-based active learning defines an efficient sampling scheme to improve the efficiency of a surrogate model of the unknown target distribution over the domain of samples.

This paper explicitly formalizes and discusses Bayesian frameworks as an active learning procedure realized through acquisition functions. In particular, pool-based active learning shows in essence a strong dualism with Bayesian frameworks. We emphasize this synergy through the dissertation on the correspondence between learning criteria and infill criteria; the former drive the sampling procedure in pool-based active learning, while the latter guide the search in Bayesian schemes through the acquisition function. This symbiosis is evidenced for the case of a single source of information adopted to query samples, and when multiple sources are at disposal of the learner to interrogate new input. Accordingly, we review and discuss popular sampling policies commonly adopted in pool-based active learning, and discern the learning criteria to accomplish a specific learning goal (Sect. 5.1). Then, the attention is dedicated to the identification of the infill criteria realized through popular acquisition functions in Bayesian optimization (Sect. 5.2). The objective is to explicitly formalize the synergy between Bayesian frameworks and active learning as

adaptive sampling schemes guided by common principles. The same avenue is followed to formalize this dualism for the case of multiple sources of information available during the learning procedure. In particular, we identify the learning criteria adopted in pool-based active learning with multiple oracles (Sect. 5.3), and compare them with the infill criteria specified by well-established multifidelity acquisition functions in multifidelity Bayesian optimization (Sect. 5.4). The objective is to clarify the shared principles and the mutual relationship that characterize the two adaptive learning schemes when the decision of the sample to query requires also the selection of the appropriate source of information to be evaluated.

5.1 Learning Criteria

Pool-based active learning determines a tailored sampling policy to ensure the maximum computational efficiency of the adaptive sampling procedure—limited and well selected amount of samples to query. This adaptive learning demands for principled guidelines to decide whether or not evaluate a certain sample based on a measure of its goodness. Learning criteria permit to establish a metric for quantifying the gains of all the possible learner decisions, and prescribe an optimal decision based in information acquired from the surrogate model. The vast majority of the literature concerning pool-based active learning identifies three essential learning criteria: informativeness, representativeness and diversity [46, 93, 126, 141, 144, 153]:

1. *Informativeness* measures the amount of information encoded by a certain sample. This means that the sampling policy is driven by the maximum likely contribution of queries that would significantly benefit the objective of the learning procedure.
2. *Representativeness* quantifies the similarity of a sample or a group of samples with respect to a target sample representative of the target distribution. Thus, the sam-

pling policy exploits the structure underlying the domain to direct the queries in locations where a sample can represent a large amount of neighbouring samples.

3. *Diversity* estimates how well the queries are disseminated over the domain of samples. This is reflected in a sampling policy that selects samples scattering across the full domain, and prevents the concentration of queries in small local regions.

Figure 3 illustrates a watering optimization problem that attempts to clarify the peculiarities of each learning criteria. This simple toy problem requires identifying the areas of a wheat field where the crop is ripe and where it is still unripe for irrigation purposes. The learning goal is formalized as the identification of the area where the wheat is lower, which means an unripe cultivation and maximum requirements for irrigation. We assume that the learner can explore a maximum of five sites on the field during the procedure. A learner driven by the pure informativeness criterion (Fig. 3a) would uniquely sample the regions of the wheat field that are likely to provide the maximum amount of information to accomplish the given learning goal; Accordingly, observations are placed where the height of the wheat is minimum and the demand for water is maximum: this maximizes the information on where it is strictly necessary to irrigate, but nothing is known about the regions where the wheat is higher and irrigation is not a priority. Conversely, a purely representative sampling (Fig. 3b) would probe the field by agglomerating observations to ensure the representativeness of the samples. This allows to partially know even areas where copious irrigation is not necessary, but increases the overall uncertainty given the small amount of samples for each agglomeration. If the learner pursues only the diversity of queries (Fig. 3c), samples would scatter the field minimizing the maximum distance between measurements. Although this allows the queries to be distributed across the entire domain, the uncertainty is high as only one sample covers a respective area of the field.




Fig. 3 Learning criteria: watering optimization problem

The remaining of this section is dedicated to the revision and discussion of popular pool-based active learning schemes. We aim to provide a broad spectrum of approaches that exemplify the implementation of different learning criteria both individually and in combination. This permits to highlight the driving principles of learning procedures, and will help to better clarify the existing synergy between active learning and Bayesian optimization accounted in the following sections. Figure 4 summarizes the relationship between the methodologies reviewed in the following and the three learning criteria.

5.1.1 Informativeness-Based

Learning procedures characterized by a pure informative criterion can be traced in uncertainty-based sampling policies. These approaches make the query decision based on the predictive uncertainty of the surrogate model, and seek to improve the density of samples in regions that exhibit the largest uncertainty with respect to a specific learning goal. Popular uncertainty-based active learning algorithms are uncertainty sampling and query-by-committee methods. Uncertainty sampling algorithms probe the domain to improve the overall accuracy of the surrogate model according to a measure of the predictive uncertainty. Examples include the quantification of the uncertainty associated with samples [74], and its alternatives as margin-based [5], least confident [77] and entropy-based [49] approaches. Other strategies defines sampling policies which promotes the minimization of the surrogate model predicted variance [17] to maximize, respectively, the decrease of loss augmenting the training set [120], and the gradient descend [13]. Other uncertainty-based strategies are query-by-committee sampling schemes [12, 155], where the most informative sample to query is selected through the maximization of the disagreement between the predictions of a committee of surrogate models computed on subsets of the locations.




Fig. 4 Mapping methodologies to learning criteria

5.1.2 Representativeness/Diversity-Based

Other pool-based active learning algorithms relies exclusively on representativeness and diversity learning frames: usually these learning criteria are implemented at the once in the learning procedure to drive the domain probing. This blend is justified by the mutual complementary relationship between representativeness and diversity: pure representativeness might concentrate the sampling in congregated representative domain regions without a proper dispersion of queries, while pure diversity might lead to the over-query of the domain and divert the learning procedure from the actual goals. The combination of both the learning criteria permits on one hand to leverage the representativeness of samples to accomplish a certain learning goal, on the other hand prevents the selection of redundant samples and high densities of queries only in circumstanced regions of the domain. Representative/diversity-based algorithms include a multitude of approaches that are commonly classified in two main schemes: clustering methodology and optimal experimental design. The former clustering algorithms identifies the most representative locations exploiting the underlying structures of the domain: the utility of samples is obtained as a function of their distance from the cluster centers. Popular examples include hierarchical clustering and k-center clustering. The former identifies a hierarchy of clusters based on the encoded information, and selects samples closer to the cluster centers [18]; the latter determines a subset of k congruent clusters that together cover the sampling space and whose radius is minimized, and the best sample minimizes the maximum distance of any point to a center [117]. The latter optimal experimental design defines a sampling policy based on a transductive approach: the learning procedure conducts the queries through a data reconstruction framework that measure the samples representativeness based on the capacity to reconstruct the training dataset. The selection of the most representative sample comes from an optimization process that maximizes the local acquisition of information about the parameters of the surrogate model [15, 35, 112].

5.1.3 Hybrid

Recent avenues explore the combination of both informativeness and representativeness/diversity learning criteria to combine the goal oriented query of the first, and the use of underlying structures preventing over-density of the second. Accordingly, combined-based algorithms integrate multiple learning criteria to improve the overall sampling performance. Those approaches are commonly classified into three main classes [153, 157]: serial-form, criteria selection, and parallel-form approaches. Serial-form algorithms use a switching approach to take advantages from all the

three learning criteria: informativeness-based techniques are used to select a subset of highly informative samples, and then representativeness/diversity techniques identify the centers of the clusters on this subset as the querying locations [126]. Criteria selection algorithms rely on a selection parameter informed by a measure of the learning improvement that suggests the appropriate learning criteria to be used during the procedure [50]. Both serial-form and criteria selection strategies combine the three learning criteria through a sequential approach where each criteria is used consecutively during the learning procedure. Parallel-form methods combine simultaneously multiple learning criteria: the utility of each sample is judged by weighting informativeness and representativeness/diversity at the same time; then, valuable samples are selected through a multi-objective optimization of the weights to maximize at the same time the improvement in terms of learning goals and the exploitation of potentially useful structures of the domain [76, 131, 136].

5.2 Acquisition Functions and Infill Criteria

The synergy between active learning and Bayesian optimization relies on the substantial analogy between the learning criteria driving the active learning procedure and the infill criteria that characterize the Bayesian learning scheme. Infill criteria provide a measure of the information gain in terms of utility acquired evaluating a certain location of the domain. In Bayesian optimization, the acquisition function is formalized according to a certain infill criterion: this permits to quantify the merit of each sample with respect to a specific learning goal. Accordingly, the sample that maximizes the querying utility is observed to enrich the learning procedure towards this goal.

In particular, Bayesian learning schemes relies on two main infill criteria: global exploration and local exploitation toward the optimum. The former exploration criterion concentrates the samples in regions of the domain where the uncertainty predicted by the surrogate is higher; this enhances the global awareness about the distribution of the objective function over the domain, but the resources might not be directed toward the goal of the procedure—e.g. minimum of the objective function. The latter exploitation criterion condenses the samples on regions where the surrogate model indicates that the objective is likely to be located—e.g. minimum of the Gaussian process mean function; exploitation realizes a goal-oriented sampling procedure that privileges the search for the objective without a potentially accurate knowledge of the overall distribution of interest. The dilemma between exploration and exploitation represents a key challenge to be carefully addressed. On one hand, a learning procedure based on pure exploration might use a large amount of samples to improve the overall

accuracy of the surrogate model without searching toward the learning goal. On the other hand, an exploitation-based learner might anchor a high density of samples to a suboptimal local solution as a consequence of information from an unreliable surrogate model. These extreme behaviours demonstrate the need to find a compromise between exploration and exploitation criteria.

In principle, infill criteria in Bayesian optimization are strongly related to the learning criteria commonly adopted in active learning. In particular:

- The concept of **exploration** is close to the **representativeness/diversity** criterion: both this learning schemes leverage underlying structures of the target distribution predicted by an accurate surrogate model to improve the awareness about the objective over the domain.
- The concept of **exploitation** is close to the **informativeness** criterion: the learner directs the selection of samples toward the believed objective without considering the global behaviour of the objective over the domain.

Figure 5 summarizes the mapping between infill criteria and learning criteria. The following sections discuss the formalization of (infill) active learning criteria for three most popular formulations of Bayesian acquisition functions, namely the expected improvement (Sect. 5.2.1), probability of improvement (Sect. 5.2.2), and max-value entropy search (Sect. 5.2.3).

5.2.1 Expected Improvement

The Expected Improvement (EI) acquisition function quantifies the expected value of the improvement in the solution of the optimization problem achieved evaluating a certain location of the domain [61, 91]. EI at the generic location x relies on the predicted improvement over the best solution of the optimization problem observed so far. Considering the




Fig. 5 Mapping of the learning criteria in active learning and infill criteria in Bayesian optimization

Gaussian process as the surrogate model for Bayesian optimization, EI can be expressed as follows:

$$U_{EI}(\mathbf{x}) = \sigma(\mathbf{x})(I(\mathbf{x})\Phi(I(\mathbf{x}))) + \mathcal{N}(I(\mathbf{x});0, 1) \quad (12)$$

where $I(\mathbf{x}) = (f(\hat{\mathbf{x}}^*) - \mu(\mathbf{x}))/\sigma(\mathbf{x})$ is the predicted improvement, $\hat{\mathbf{x}}^*$ is the current location of the best value of the objective sampled so far, $\Phi(\cdot)$ is the cumulative distribution function of a standard normal distribution, μ is the mean function and σ is the standard deviation of the GP. The computation of $U_{EI}(\mathbf{x})$ requires limited computational resources and the first-order derivatives are easy to calculate:

$$\frac{\partial U_{EI}(\mathbf{x})}{\mu(\mathbf{x})} = -\Phi(I(\mathbf{x})) \quad (13)$$

$$\frac{\partial U_{EI}(\mathbf{x})}{\sigma(\mathbf{x})} = \phi(I(\mathbf{x})). \quad (14)$$

Both Equations (13) and (14) demonstrate that $U_{EI}(\mathbf{x})$ is monotonic with respect to the increase of both the mean and the uncertainty of the GP surrogate model. This highlights a form of trade-off between exploration and exploitation: the formulation of the EI permits to balance the sampling in locations of the domain where is likely to have a significant improvement of the solution with respect to the current best solution, and the observations of regions where the improvement might be contained but the prediction is highly uncertain. In principle, it is possible to state that EI is driven by a combination of informativeness and representativeness/diversity criteria adopted in active learning. On one hand, the learner seeks to direct the computational resources towards the maximization of the learning contribution and achievement of the goal—informativeness; on the other hand, the learner pursues the awareness of the objective distribution over the domain to improve the quality of the prediction and better drive the search—representativeness/diversity. The predictive framework of the surrogate model regulates the learning thrusts privileging the one over the other on the basis of the information about the objective function acquired over the iterations.

5.2.2 Probability of Improvement

The Probability of Improvement (PI) acquisition function targets the locations characterized by the highest probability of achieving the goal, based on the information from the current surrogate model [60, 69]. PI measures the probability that the prediction of the surrogate model at the generic location is lower than the best observation of the objective function so far. Under the Gaussian process surrogate model, the PI acquisition function is computed in closed form as follows:

$$U_{PI}(\mathbf{x}) = \Phi(I(\mathbf{x})) \quad (15)$$

where $\Phi(\cdot)$ is the cumulative distribution function of a standard normal distribution and \mathbf{x}^* is the current location of the best value of the objective. Similarly to EI, also $U_{PI}(\mathbf{x})$ is inexpensive to compute and the evaluation of the first-order derivatives requires simple calculations:

$$\frac{\partial U_{PI}(\mathbf{x})}{\partial \mu(\mathbf{x})} = -\frac{1}{\sigma(\mathbf{x})}\phi(I(\mathbf{x})) \quad (16)$$

$$\frac{\partial U_{PI}(\mathbf{x})}{\partial \sigma(\mathbf{x})} = -\frac{I(\mathbf{x})}{\sigma(\mathbf{x})}\phi(I(\mathbf{x})) \quad (17)$$

where ϕ is the standard Gaussian probability density function. As demonstrated by Equation (16), regions of the input space characterized by lower values of the posterior mean of the GP are preferred for sampling, at fixed uncertainty of the surrogate. Moreover, Equation (17) shows that if $\mu(\mathbf{x}) < f(\mathbf{x}^*)$ the regions characterized by lower uncertainty are preferred and, conversely, PI increases with uncertainty. Overall, the PI acquisition function can be considered as an exploitative scheme that determines the most informative location as the one that potentially produces a larger reduction of the minimum value of the objective function observed so far. This is achieved sampling regions where the surrogate model is reliable and characterized by lower levels of uncertainty. In principle, this sampling scheme makes PI in accordance with the informativeness criterion: the search toward the optimum is uniquely directed in regions of the domain that exhibit the higher probability of achieving the goal according to the emulator prediction.

5.2.3 Entropy Search and Max-Value Entropy Search

The Entropy Search (ES) acquisition function measures the differential entropy of the believed global minimum location of the objective function, and targets the reduction of uncertainty selecting the sample that maximizes the decrease of differential entropy [47]. The ES acquisition function is formulated as follows:

$$U_{ES}(\mathbf{x}) = H(p(x^*|\mathcal{D})) - \mathbb{E}_{f(\mathbf{x})|\mathcal{D}}[H(p(x^*|f(\mathbf{x}), \mathcal{D}))] \quad (18)$$

where $H(p(x^*))$ is the entropy of the posterior distribution at the current iteration on the location of the minimum of the objective function x^* , and $\mathbb{E}_{f(\mathbf{x})}[\cdot]$ is the expectation over $f(\mathbf{x})$ of the entropy of the posterior distribution at the next iteration on x^* . Typically, the exact calculation of the second term of Equation (18) is not possible and requires complex and expensive computational techniques to provide an approximation of $U_{ES}(\mathbf{x})$.

The Max-value entropy search (MES) [135] acquisition function is derived from the ES acquisition function and

allows to reduce the computational effort required to estimate Equation (18) measuring the differential entropy of the minimum-value of the objective function:

$$U_{MES}(\mathbf{x}) = H(p(f|\mathcal{D})) - \mathbb{E}_{f(\mathbf{x})|\mathcal{D}}[H(p(f|f^*, \mathcal{D}))] \quad (19)$$

where the first and the second term are now computed on the minimum value of the objective function f^* . This permits to simplify the computations and to approximate the second term through a Monte Carlo strategy [135]. The analysis of the derivatives is not possible for the MES acquisition function since the formulation of the second term of Equation (19) is intractable.

As reported by Wang et al. [135] in their experimental analysis, MES targets the balance between the exploration of locations characterized by higher uncertainty of the surrogate model, and the exploitation toward the believed optimum of the objective function. However, Nguyen et al. [98] demonstrate that MES might suffer from an imbalanced exploration/exploitation trade-off due to noisy observations of the objective function, and to the discrepancy in the computation of the mutual information in the second term of Equation (19). As a result, MES might over-exploit the domain in presence of noise in measurements, and over-explore when the discrepancy in the evaluation issue determines a pronounced sensitivity to the uncertainty of the surrogate model. Overall, the adaptive sampling scheme determined by the MES acquisition function follows both the informativeness and the representativeness/diversity learning criteria: the most promising sample is ideally selected targeting the balance between the search toward the believed minimum predicted by the emulator, and the decrease of uncertainty about the objective function distribution.

5.3 Learning Criteria with Multiple Oracles

Most of the active learning paradigms rely on a unique and supposed omniscient source of information about the target distribution. This oracle is iteratively queried by the learner to evaluate the value of the distribution at certain locations, and is assumed that its estimate is exact. In many other scenarios, the learner can elicit information from multiple imperfect oracles at different levels of reliability, accuracy and cost. Accordingly, the active learning community introduces a multitude of annotator-aware algorithms which are capable to efficiently learn from multiple sources of information. This requires to make an additional decision during the learning procedure: the learner has to select at each iteration the most useful sample and the associated information source to query. In this context, the original learning criteria of informativeness and representativeness/diversity (Sect. 5.1) evolve and extend to quantify the utility

of querying the domain with a certain level of accuracy and associated cost:

1. *Informativeness* seeks to maximize the amount of information from deciding the sample and information source to query. Thus, the learner might privilege the evaluations from accurate and yet costly oracles to capitalize from high-quality information and potentially reach the objective.
2. *Representativeness* attempts to identify underlying structures of the domain to better inform the search procedure. In this case, the decision making process might prefer to interrogate less expensive sources of information to contain the required effort, especially if cheap predictions of the target distribution exhibit good correlation with the estimate of the accurate oracle.
3. *Diversity* scatters the sampling effort over the domain to pursue a proper distribution of evaluations and augment the awareness about the target distribution. This might be favored by a major use of less accurate predictions of the target distribution, which are more likely to well address the cost/effectiveness trade-off during the diversity sampling.

The remaining of this section provides an overview of different multiple oracles active learning methodologies to present and further clarify popular extensions of the learning criteria to a multi-oracle setting.

Typically, active learning paradigms are extended to the multiple-oracle setting through relabeling, repeating-labeling, probabilistic and transfer knowledge, and cost-aware algorithms. Relabeling approaches query samples multiple times using the library of sources of information available, and the final query is obtained via majority voting [156]. Popular methodologies following this scheme pursue the identification of a subset of oracles according to the proximity of their upper confidence bound to the maximum upper confidence bound, and apply the majority voting technique only considering the queries of this informative subset [25]. Other multi-oracle active learning methods use a repeating-labeling procedure: the learner integrates the repeated—often noisy—prediction of the oracles to improve the quality of the evaluation process and the accuracy of the surrogate model learned from data [54]. Both relabeling and repeating-labeling approaches share a common drawback: the same unknown sample is evaluated multiple times with different oracles, which results in a sub-optimal usage of the available sources of information. Probabilistic and transfer learning methodologies attempt to overcome this limitation. Probabilistic frameworks rely on surrogate models specifically conceived for the multi-source scenario that provide a predictive framework to estimate the accuracy of each oracle in the evaluation of samples over the domain [148, 149].

Transfer knowledge approaches enhance the simultaneous selection of the most informative location to sample and the associated most profitable source to query; this is achieved through the transfer of knowledge from samples not evaluated in auxiliary domains to support the estimate of the oracle reliability [31]. Recent advancements in multiple oracles active learning are cost-effective algorithms, where the cost of an oracle is evaluated considering both the overall reliability of the prediction and the quality of samples in specific locations [36, 52, 151]. The cost-effectiveness property enhance the use of computational resources for the evaluation of samples, and targets the search toward the learning objectives while guarantees an optimal trade-off between evaluation accuracy and computational cost.

From the examined literature, the three learning criteria appear frequently coupled together during the learning procedure with multiple sources to query. This appears as a natural evolution of what has already been observed in the literature for active learning with single information source: the overall learning procedure usually benefits from a balanced learning scheme driven by informativeness and representativeness/diversity. In particular, informativeness permits to direct the search toward the learning goal, while representativeness/diversity augments the learner awareness about the target distribution over the domain; the combination of these learning criteria—in different measures—contributes to improve the performance of the active learning algorithms by using efficiently the computational resources and the information from multiple oracles.

5.4 Multifidelity Acquisition Functions and Infill Criteria

This section further investigates and highlights the synergy between active learning and Bayesian optimization for the specific case of multiple source of information used to accomplish the learning goal. Similarly to the single source setting, this symbiotic relationship is revealed through common principles characterizing the infill criteria in multifidelity Bayesian optimization and the learning criteria in active learning with multiple oracles. The multifidelity scenario imposes an additional decision to be made: the learner has to identify the appropriate information source to query according to an accuracy/cost trade-off. This is reflected in the formalization of infill criteria capable to define an efficient and balanced sampling policy, targeting either the wise selection of the samples and the level of fidelity which ensures the maximum benefits with the minimum cost. Accordingly, the multifidelity acquisition function formalizes an adaptive sampling scheme based on one or multiple infill criteria to quantify the utility of querying a location of the domain with a specific level of fidelity.

Based on this considerations, the exploration and exploitation infill strategies are extended according to the peculiarities of the multifidelity setting:

- *Exploration* is close to the representativeness/diversity criterion and defines a sampling policy that incentivizes the overall reduction of the surrogate uncertainty. Accordingly, the selection of the appropriate level of fidelity is driven by a trade-off between accuracy and evaluation cost. This might be accomplished through less-expensive low-fidelity information to contain the demand for computational resources during exploration.
- *Exploitation* is close to the informativeness criterion: concentrates the sampling process in the regions of the domain where optimal solutions are likely to be located. For this purpose, the learner might emphasize the use of accurate evaluations of the target function to refine the solution of the learning procedure toward the specific goal.

Similarly to the acquisition functions in Bayesian optimization (Sect. 5.2), the symmetry between informativeness and exploitation criterion, and between representativeness/diversity and exploration criterion is preserved in the multifidelity setting. The following sections are dedicated to the revision and discussion of popular multifidelity acquisition function, namely the multifidelity expected improvement (Sect. 5.4.1), multifidelity probability of improvement (Sect. 5.4.2) and multifidelity max-value entropy search (Sect. 5.4.3). The goal is to highlight the equivalent principles driving both the learning schemes, and further clarify the elements that encode the symbiotic relationship that exists between multifidelity Bayesian optimization and multi-oracle active learning.

5.4.1 Multifidelity Expected Improvement

The Multifidelity Expected Improvement (MFEI) extends the expected improvement acquisition function to define a learning scheme in the multifidelity setting as follows [51]:

$$U_{MFEI}(\mathbf{x}, l) = U_{EI}(\mathbf{x}, L)\alpha_1(\mathbf{x}, l)\alpha_2(\mathbf{x}, l)\alpha_3(\mathbf{x}, l) \quad (20)$$

where $U_{EI}(\mathbf{x}, L)$ is the expected improvement depicted in Equation (12) evaluated at the highest level of fidelity L , and the utility functions α_1 , α_2 and α_3 are defined as follows:

$$\alpha_1(\mathbf{x}, l) = \text{corr}[f^{(l)}, f^{(L)}] \quad (21)$$

$$\alpha_2(\mathbf{x}, l) = 1 - \frac{\sigma_\epsilon}{\sqrt{\sigma^{2(l)}(\mathbf{x}) + \sigma_\epsilon^2}} \quad (22)$$

$$\alpha_3(l) = \frac{\lambda^{(L)}}{\lambda^{(l)}}. \quad (23)$$

The first element α_1 is the posterior correlation coefficient between the level of fidelity l and the high-fidelity level L , and accounts for reduction of the expected improvement when a sample is evaluated with a low fidelity model. This term reflects a measure of the informativeness of the l -th source of information at the location \mathbf{x} , and balances the amount of improvement achievable evaluating the high-fidelity level L with the reliability of the prediction associated with the level of fidelity l . Accordingly, α_1 modifies the learning scheme by adding a penalty in the formulation that reduces the U_{MFEI} when $1 \leq l < L$: this includes awareness about the increase of uncertainty associated with a low-fidelity prediction. The second element α_2 is conceived to adjust the expected improvement when the output at the l -th level of fidelity contains random errors. This is equivalent to consider the reduction of the uncertainty on the Gaussian process prediction after a new evaluation of the objective function is added to the dataset \mathcal{D} . This function allows to improve the robustness of U_{MFEI} when the representation of $f^{(l)}$ at different levels of fidelity is affected by noise in the measurements. The third element α_3 is formulated as the ratio between the computational cost of the high-fidelity level L and the l -th level of fidelity. This permits to balance the informative contributions of high- and a lower-fidelity observation and the related computational resources required for the evaluation. The effect of this term is to encourage the use of low-fidelity representations if almost the same expected improvement can be achieved with a high-fidelity evaluation. This directs wisely the use of computational resources to achieve the representativeness/diversity of samples, and prevents a massive use of expensive accurate queries during exploration phases.

5.4.2 Multifidelity Probability of Improvement

The Multifidelity Probability of Improvement (MFPI) acquisition function provides an extended formulation of the probability of improvement suitable for the multifidelity scenario as follows [114]:

$$U_{MFPI}(\mathbf{x}, l) = U_{PI}(\mathbf{x}, L)\eta_1(\mathbf{x}, l)\eta_2(l)\eta_3(\mathbf{x}, l) \quad (24)$$

where the PI acquisition function (Equation (15)) is computed considering the highest-fidelity level L available, and the utility function η_1 , η_2 and η_3 are defined as follows:

$$\eta_1(\mathbf{x}, l) = \text{corr}[f^{(l)}, f^{(L)}] \quad (25)$$

$$\eta_2(l) = \frac{\lambda^{(L)}}{\lambda^{(l)}} \quad (26)$$

$$\eta_3(\mathbf{x}, l) = \prod_{i=1}^{n_l} [1 - R(\mathbf{x}, \mathbf{x}_i^{(l)})]. \quad (27)$$

The first term η_1 shares the same formalization of the utility function α_1 in Equation (21), and accounts for the increase of uncertainty associated with low-fidelity representations $1 \leq l < L$ if compared with the high-fidelity output L . This reduces the probability of improvement if a low-fidelity representation is queried at a specific location of the input space \mathbf{x} . As already highlighted in Sect. 5.4.1, η_1 incentivates a form of informativeness learning where the information source is selected according to its capability to accurately represent the objective function. Similarly, the second utility function η_2 is also included in the multifidelity expected improvement in Equation (23) as the α_3 term. This element balances the computational costs and the informative contributions achieved through the l -th level of fidelity. This prevents the rise of computational demand produced by the over-exploitative nature of the probability of improvement (Sect. 5.2.2): η_2 encourages the use of fast low-fidelity data if the discrepancy between the l -th level of fidelity and the high-fidelity L —quantified by η_1 —is not significant. The third element η_3 is the sample density function and is computed as the product of the complement to unity of the spatial correlation function $R(\cdot)$ [80] evaluated for the n_l samples considering the l -th level of fidelity. This term reduces the probability of improvement in locations with an high sampling density—over exploitation of the domain—to prevent the clustering of data. Accordingly, η_3 promotes a form of representativeness/diversity learning scheme and encourages the exploration to augment the awareness about the domain structure.

5.4.3 Multifidelity Entropy Search and Multifidelity Max-Value Entropy Search

The Multifidelity Entropy Search (MFES) acquisition function is formulated extending the entropy search acquisition function to query multiple sources of information [154]

$$U_{MFES}(\mathbf{x}) = H(p(x^* | \mathcal{D})) - \mathbb{E}_{f^{(l)}(\mathbf{x}) | \mathcal{D}} [H(p(x^* | f^{(l)}(\mathbf{x}), \mathcal{D}))] \quad (28)$$

where the expectation term $\mathbb{E}_{f^{(l)}(\mathbf{x})}[\cdot]$ considers multiple levels of fidelity $l = 1, \dots, L$. Similarly to the entropy search acquisition function, the computation of the expectation in Equation (28) is not possible in closed-form and requires an intensive procedure to provide a reliable approximation.

The Multifidelity Max-Value Entropy Search (MFMES) acquisition function can be formulated extending the max-value entropy search to a multifidelity setting as follows [130]:

$$U_{MFMES}(\mathbf{x}) = [H(p(f^{(l)}|\mathcal{D})) - \mathbb{E}_{f^{(l)}(\mathbf{x})|\mathcal{D}}[H(p(f^{(l)}|f^{*(L)}, \mathcal{D}))]]/\lambda^{(l)} \quad (29)$$

where the differential entropy is measured on the minimum value of the objective function $f^{*(L)}$ considering the high-fidelity representation L . In this case, the approximation of the expectation term in Equation (29) relies on a Monte Carlo strategy that allows to contain the computational cost if compared with the procedure used for the MFES acquisition function [130].

In the multifidelity scenario, the FMES acquisition function measures the information gain obtained evaluating the objective function $f^{(l)}(\mathbf{x})$ at a certain location \mathbf{x} and associated level of fidelity l with respect to the global minimum of the objective function. This can be interpreted as an informativeness-driven learning based on the reduction of the uncertainty associated with the minimum value of the objective $f^{*(L)}$ through the observation $f^{(l)}(\mathbf{x})$, where this uncertainty is measured as the differential entropy associated with the l -th level of fidelity. At the same time, the information gain is also sensitive to the accuracy of the surrogate predictive framework, and realizes a form of representativeness/diversity balancing to improve the awareness about the distribution of the objective function over the domain. The sensitivity to the computational cost $\lambda^{(l)}$ of the l -th level of fidelity is introduced in Equation (29) to balance the quality of the source—quantified by the information gain—and the demand for computational resources.

6 Experiments

This section investigates and compares the performance of the acquisition functions for both single-fidelity and multifidelity Bayesian optimization over a set of benchmark problems conceived to stress the algorithms. The objective is to highlight advantages and opportunities offered by different learning principles over challenging mathematical properties of the objective function, which are frequently encountered in real-world engineering and scientific problems. [83]. In particular, this comparative study considers the expected improvement (Sect. 5.2.1), probability of improvement (PI) (Sect. 5.2.2), and Max-Value Entropy Search (MES) (Sect. 5.2.3) for the single-fidelity frameworks, and their multifidelity counterparts Multifidelity Expected Improvement (MFEI) (Sect. 5.4.1), Multifidelity Probability of Improvement (MFPI) (Sect. 5.4.2) and Multifidelity Max-Value Entropy Search (FMES) (Sect. 5.4.3).

We impose the same initialization conditions for both the single-fidelity and the multifidelity algorithms. This initial setting includes: (i) the initial dataset of $N_0^{(l)}$ samples for each

level of fidelity l to compute the prior surrogate model of the objective function, (ii) the computational cost assigned to each level of fidelity $\lambda^{(l)}$, and (iii) the maximum computational budget B_{max} allocated for each benchmark problem defined linearly with the dimensionality D of the problem $B_{max} = 100D$. The initial dataset $N_0^{(l)}$ is obtained through Latin hypercube sampling for all the numerical experiments [87] to ensure the full coverage of the range of the optimization variables. The computational budget $B = \sum \lambda_i^{(l)}$ is quantified as the cumulative computational cost used during the optimization at each iteration i .

All the methods are based on the Gaussian processes surrogate model and its extension to the multifidelity setting. We implement the square exponential kernels for all the GP covariances, and use the maximum likelihood estimation approach to optimize the hyperparameters of the kernel and the mean function of the GP [127].

6.1 Benchmark Problems

The following set of benchmark problems is specifically conceived to investigate the capabilities of different learning criteria over challenging mathematical properties of the objective function [83]. In particular, the experimental settings include a variety of attributes that can be traced in real-world optimization problems, namely local and global behaviours, nonlinearities and discontinuities, multimodality and noise. The set of problems consists of several objective functions such as the Forrester continuous and discontinuous, the Rosenbrock increasing the domain dimensionality, the Rastrigin shifted and rotated, the Agglomeration of Locally Optimized Surrogate (ALOS), a coupled spring-mass optimization problem and the noisy Paciorek function.

6.1.1 Forrester Function


The Forrester function is a popular test-case to investigate the performance of different learning strategies over a non-linear one-dimensional distribution characterized by local behaviours. This benchmark problem guarantees an high interpretability of the results thanks to the one-dimensional nature of the objective function. The search domain is bounded between $\chi = [0, 1]$ and four levels of fidelity are available during the optimization:

$$f^{(4)}(\mathbf{x}) = (6\mathbf{x} - 2)^2 \sin(12\mathbf{x} - 4) \quad (30)$$

$$f^{(3)}(\mathbf{x}) = (5.5\mathbf{x} - 2.5)^2 \sin(12\mathbf{x} - 4) \quad (31)$$

$$f^{(2)}(\mathbf{x}) = 0.75f_1(\mathbf{x}) + 5(\mathbf{x} - 0.5) - 2 \quad (32)$$

Fig. 6 Forrester function benchmark problems



$$f^{(1)}(x) = 0.5f_1(x) + 10(x - 0.5) - 5 \tag{33}$$

where $f^{(4)}$ is the high-fidelity function and the levels of fidelity $l = 1, 2, 3, 4$ increase with the accuracy of the representations. Figure 6(a) reports the four levels of fidelity for the Forrester function over the search domain. The analytical minimum of the Forrester function is equal to $f^{*(4)} = -6.0207$ and is located at the domain point $x^* = 0.7572$.

6.1.2 Jump Forrester Function

The jump Forrester function introduces a discontinuity in the formulation of the Forrester function to investigate the capabilities of learning schemes to refine the surrogate model and capture the instantaneous variation of the objective function over the domain. This scenario can often occur in real problems where the phenomena of interest—e.g. physical quantity of interest in engineering—evolves over the domain and determine large variations of the objective function values. Figure 6(b) reports the two levels of fidelity are available during the search procedure:

$$f^{(2)}(x) = \begin{cases} (6x - 2)^2 \sin(12x - 4), & 0 \leq x \leq 0.5 \\ (6x - 2)^2 \sin(12x - 4) + 10, & 0.5 < x \leq 1 \end{cases} \tag{34}$$

$$f^{(1)}(x) = \begin{cases} 0.5f^{(2)}(x) + 10(x - 0.5) - 5, & 0 \leq x \leq 0.5 \\ 0.5f^{(2)}(x) + 10(x - 0.5) - 2, & 0.5 < x \leq 1 \end{cases} \tag{35}$$


where $f^{(2)}$ is the high-fidelity information source. The optimum is located at $x^* = 0.75724876$ corresponding to a value of the objective equal to $f^{*(2)} = -0.9863$.

6.1.3 Rosenbrock Function

The Rosenbrock function permits to investigate the learning criteria over a non-convex objective function that allows for parametric scalability over the domain $\chi = [-2, 2]^D$ where D is the dimensionality of the input space. A library of three levels of fidelity is available (Fig. 7):

$$f^{(3)}(x) = \sum_{i=1}^{D-1} 100(x_{i+1} - x_i^2)^2 + (1 - x_i)^2 \tag{36}$$

Fig. 7 Rosenbrock function benchmark problem over the $D = 2$ dimensional domain



$$f^{(2)}(\mathbf{x}) = \sum_{i=1}^{D-1} 50(x_{i+1} - x_i^2)^2 + (-2 - x_i)^2 - \sum_{i=1}^D 0.5x_i \tag{37}$$

$$f^{(1)}(\mathbf{x}) = \frac{f^{(3)}(\mathbf{x}) - 4 - \sum_{i=1}^D 0.5x_i}{10 + \sum_{i=1}^D 0.25x_i} \tag{38}$$


where the high-fidelity function is $f^{(3)}$ and the lower-fidelities are obtained using a transformation of $f^{(3)}$ based on linear additive and multiplicative factors. The analytical minimum is located at $x^* = [1, 1]^D$ and corresponds to a value of the objective function $f^{*(3)} = 0$. The scalability of the Rosenbrock function with D of the formulation allows to test the performance of the methods at increasing dimensionality of the input space. In this study, we consider the cases $D = 2, 5, 10$.

6.1.4 ALOS Functions

The Agglomeration of Locally Optimized Surrogate (ALOS) is a heterogeneous and non-polynomial function defined on unit hypercubes up to three dimensions useful to assess the accuracy of surrogate models in presence of localized behaviours. In particular, the ALOS function reproduces a real-world scenario where the objective functions is characterized by oscillatory phenomena at different frequency distributed along the domain. We consider two levels of fidelity and increasing dimensionality of the input space $D = 1, 2, 3$. For $D = 1$ the ALOS function is formalized as follows:

$$\begin{cases} f^{(2)}(\mathbf{x}) = \sin[30(\mathbf{x} - 0.9)^4] \cos[2(\mathbf{x} - 0.9)] \\ \quad + (\mathbf{x} - 0.9)/2 \\ f^{(1)}(\mathbf{x}) = (f^{(2)}(\mathbf{x}) - 1.0 + \mathbf{x}) / (1.0 + 0.25\mathbf{x}) \end{cases} \tag{39}$$

Fig. 8 ALOS function benchmark problems over the $D = 1$ and $D = 2$ dimensional domain



(a) ALOS $D = 1$

and for $D = 2, 3$ is formulated as:

$$\begin{cases} f^{(2)}(\mathbf{x}) = \sin[21(x_1 - 0.9)^4] \cos[2(x_1 - 0.9)] \\ \quad + (x_1 - 0.7)/2 + \sum_{i=2}^D ix_i \sin\left(\prod_{j=1}^i x_j\right) \\ f^{(1)}(\mathbf{x}) = (f^{(2)}(\mathbf{x}) - 2.0 + \sum_{i=1}^D x_i) / (5.0 \\ \quad + \sum_{i=1}^2 0.25ix_i - \sum_{i=3}^D 0.25ix_i) \end{cases} \tag{40}$$

For $D = 1$, the analytical optimum is located at $x^* = 0.2755$ corresponding to $f^{*(2)} = -0.6250$ while for $D \geq 2$ the minimum is located at $x^* = [0, 0]^D$ with value of the objective function $f^{*(2)} = -0.5627123$. Figure 8 illustrates the high and low-fidelity ALOS function for $D = 1$ (Fig. 8a) and $D = 2$ (Fig. 8b).

6.1.5 Shifted-Rotated Rastrigin Function


The Rastrigin function is commonly used as test function to represent real-world applications where the objective function might present an high multimodal behaviour. We adopt a benchmark problem based on the original formulation of the Rastrigin function shifted and rotated as follows (Fig. 9):

$$f(\mathbf{z}) = \sum_{i=1}^D (z_i^2 + 1 - \cos(10\pi z_i)), \tag{41}$$

where: $\mathbf{z} = R(\theta)(\mathbf{x} - x^*)$ and $R(\theta) = \begin{bmatrix} \cos \theta & -\sin \theta \\ \sin \theta & \cos \theta \end{bmatrix}$ is the rotation matrix with the rotation angle fixed at $\theta = 0.2$. We define three levels of fidelity for this benchmark problems as follows:


$$f^{(l)}(\mathbf{z}, \phi) = f(\mathbf{z}) + e_r(\mathbf{z}, \phi_i) \tag{42}$$

where $e_r(\mathbf{z}, \phi_i)$ is the resolution error:



(b) ALOS $D = 2$

Fig. 9 Rastrigin function shifted and rotated benchmark problem



$$e_r(\mathbf{z}, \phi) = \sum_{i=1}^2 a(\phi) \cos^2(w(\phi)z_i + b(\phi) + \pi). \tag{43}$$

with $\Theta(\phi) = 1 - 0.0001\phi$, $a(\phi) = \Theta(\phi)$, $w(\phi) = 10\pi\Theta(\phi)$, and $b(\phi) = 0.5\pi\Theta(\phi)$. Thus, we define the high-fidelity function $f^{(3)}(\phi = 10000)$, the intermediate fidelity function $f^{(2)}(\phi = 5000)$ and the low-fidelity function $f^{(1)}(\phi = 2500)$. For this benchmark, the input variables are defined within the interval $\chi = [-0.1, 0.2]^2$ and the analytical optimum is $f^{*(3)} = 0$ located at $x^* = [0.1, 0.1]$.

6.1.6 Spring-Mass System

This benchmark problem consists of a coupled spring mass system composed of two masses connected by two springs. The challenges associated with this simple physical optimization problem are related to the intrinsic multimodality induced by the elastic behaviour of the system dynamics. We consider the masses m_1 and m_2 concentrated at their center of gravity and the elastic behaviour of the two spring modeled through the Hooke’s law and characterized by the Hooke’s constants k_1 and k_2 , respectively. Considering a friction-less dynamics, it is possible to define the equations of motion as follows

$$m_1 \ddot{h}_1(t) = (-k_1 - k_2)h_1(t) + k_2h_2(t) \tag{44}$$

$$m_2 \ddot{h}_2(t) = k_2h_1(t) + (-k_1 - k_2)h_2(t). \tag{45}$$

where $h_1(t)$ and $h_2(t)$ are the positions of the masses as a function of time t .

Equation (44) can be solved using the fourth-order accurate Runge–Kutta time-marching method and varying the time-step dt to define two fidelity levels. Specifically, we define the high-fidelity model $f^{(2)}(dt = 0.01)$ and the low-fidelity model $f^{(1)}(dt = 0.6)$. The benchmark problem consists in the identification of the combination of masses and Hooke’s constants of spring $\mathbf{x} = [m_1, m_2, k_1, k_2]$ that minimizes $h_1(t = 6)$ considering the domain $\chi = [1, 4]^4$ and the initial conditions of motion $h_1 = h_2 = 0$ and $\dot{h}_1 = \dot{h}_2 = 0$.

6.1.7 Paciorek Function with Noise

The Paciorek function reproduces an optimization settings where the objective function is affected by measurement noise and localized multimodal behaviour. This scenario is replicated through a uniformly distributed random noise over the two levels of fidelity defined as follows (Fig. 10):

$$f^{(2)}(\mathbf{x}) = \sin\left(\prod_{i=1}^D x_i\right)^{-1} \tag{46}$$

$$f^{(1)}(\mathbf{x}) = f^{(2)}(\mathbf{x}) - 9A^2 \cos\left(\prod_{i=1}^D x_i\right)^{-1} + \text{rand.norm}(0, \alpha) \tag{47}$$

where $A = 0.5$, $\alpha = 0.2$, and the input variable is defined across the input domain $\chi = [0, 3, 1]^2$.

6.2 Results and Discussion

First, we define the following evaluation metrics to assess the performances of the Bayesian schemes [83]:




Fig. 10 Paciorek function benchmark problem

$$\epsilon_x = \frac{\|\mathbf{x}^* - \hat{\mathbf{x}}^*\|}{\sqrt{N}} \quad (48)$$

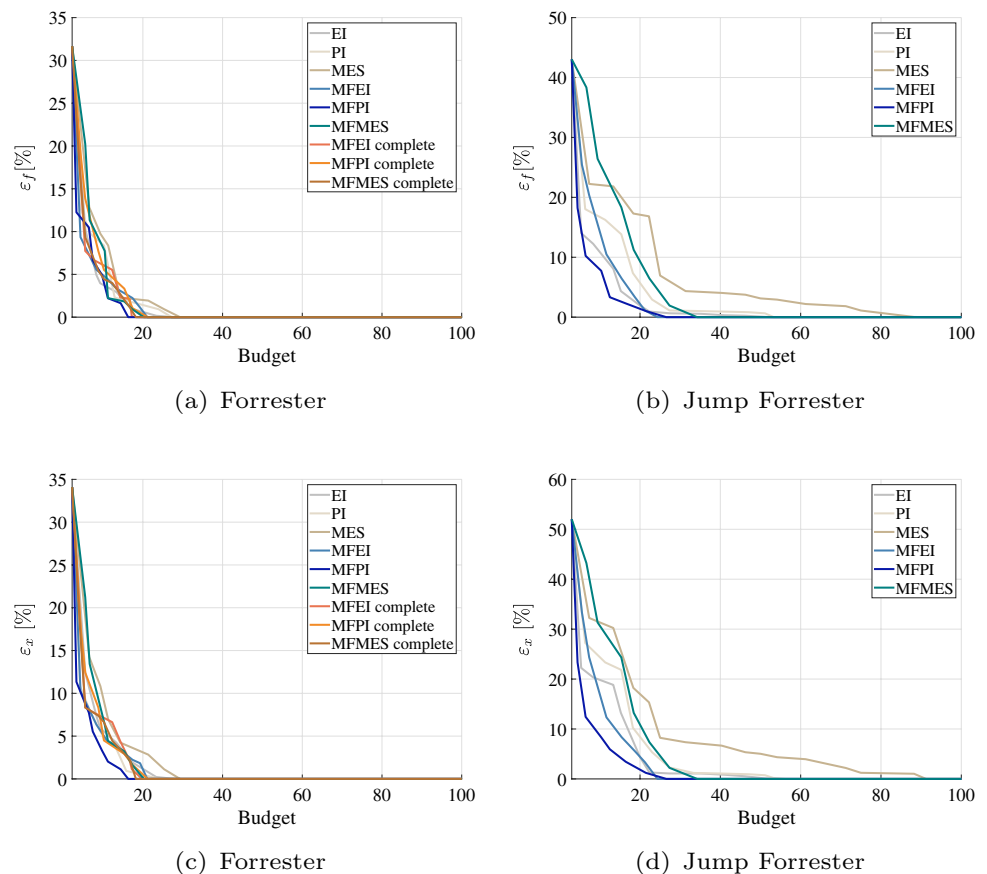
$$\epsilon_f = \frac{f(\hat{\mathbf{x}}^*) - f^*}{f_{\max} - f^*} \quad (49)$$

where \mathbf{x}^* is the location of the analytical optimum, $\hat{\mathbf{x}}^*$ is the optimum identified by the algorithm, and f_{\max} and f^* are the maximum and minimum of the objective function, respectively. The first metric ϵ_x quantifies the search error in the domain of the objective function, while the second metric ϵ_f evaluates the error associated with the learning goal—minimum of the objective function [118]. We evaluate the metrics ϵ_x and ϵ_f as functions of the computational budget B defined as the cumulative computational cost associated with observations of the objective function at the l -th level of fidelity. We run 10 trails for each benchmark problem presented in Sect. 6.1 to compensate the influence of the random initial design of experiments, and to verify the sensitivity and robustness of the algorithms to the initialization setting. The results for all the experiments are reported in terms of median values of ϵ_x and ϵ_f .

Figure 11 summarizes the outcome obtained for the Forrester function and discontinuous Forrester function. The

results for the Forrester benchmark (Fig. 11a, c) show that the multifidelity algorithms identify the optimum solution with a significant reduction of the computational budget if compared with the single fidelity counterparts. The best performing algorithm is the MFPI learner considering only the high-fidelity $l = 4$ and the lower-fidelity $l = 1$ levels, while the second best is the MFEI acquisition function considering available the complete spectrum of fidelities $l = 1, 2, 3, 4$. These outcomes suggest that multifidelity learning paradigms driven majorly by informativeness—MFPI acquisition function—are capable to efficiently direct the computational resources toward the optimum of low-dimensional objective functions in presence of continuous localized behaviour. Moreover, it should be noted that the MFEI capitalizes from all the information sources available and leverages the balance between informativeness and representativeness/diversity to effectively search toward the analytical optimum. The single fidelity Bayesian frameworks exhibit a lower convergence rate with respect to the multifidelity algorithms. The EI and PI uses almost the same computational budget to identify the optimum solution, while the MES adopts more evaluations of the objective function. This confirms the observations for the multifidelity experiments. PI takes advantage from the purely exploitation of high-fidelity samples in the surrounding of the surrogate minimum to reach

Fig. 11 Performances of the competing algorithms for the Forrester and Jump Forrester benchmarks



the optimum. This can be explained with the computation of an accurate surrogate model—at least close to the optimum—for low-dimensional objective functions. In contrast, EI balances an exploration phase to improve the overall accuracy of the surrogate with the exploitation toward the believed optimum. Particular attention should be dedicated to the MES and MFMES outcomes. In the single-fidelity frameworks, MES scores slightly worst both in terms of convergence rate and budget expenditure. This can be interpreted with an overall over-exploration behavior: MES distributes computational resources to explore the domain and refine the surrogate model, and directs lately efforts toward the optimum. This trend is considerably dampened in the multifidelity scenario, where MFMES shows good capabilities especially when all the sources of information are available during the search. In this case, cheap low-fidelity observations are used to explore the domain with contained computational expenditure, and high-fidelity data are mostly adopted to search toward the prescribed optimum location.

The discontinuous Forrester problem introduce a discontinuous local property of the objective function that further stresses the learning schemes. This can be explicitly observed with the average improvement of the budget required to achieve the optimum. Overall, it is possible to identify the same trends observed for the continuous Forrester function (Fig. 11b, d): either balancing exploration

and exploitation—EI and MFEI—or a major exploitation search—PI and MFPI—lead to an efficient identification of the analytical optimum. In contrast, the over-exploration of MES and MFMES decelerates the optimization procedure with respect to the counterpart competing methods. This can be observed majorly for the MES which uses almost all the budget available to explore the domain and finally reach the optimum.

Figure 12 illustrates the experiments conducted on the Rosenbrock benchmark function increasing the dimensionality D of the domain. This allows to investigate the performance of the learning scheme as the number of parameters to optimize increases. Overall, the multifidelity schemes deliver better convergences with a fraction of the computational budget required by single-fidelity algorithms for all the dimensions of the domain— $D = 2, 5, 10$.

For $D = 2$ (Fig. 12a, b), MFEI and MFPI implementing only the highest and lower levels of fidelity $l = 1, 3$ are the best performing algorithms, followed by the counterpart considering all the fidelities spectrum and the MFMES also learning from $l = 1, 3$. Two major observations can be made in this experimental setting. First, multifidelity learners are not capable to make advantage of the intermediate fidelity $l = 2$ during exploration leading to an increase of the computational expenditure. A possible explanation to these outcomes is the local behaviour of the intermediate

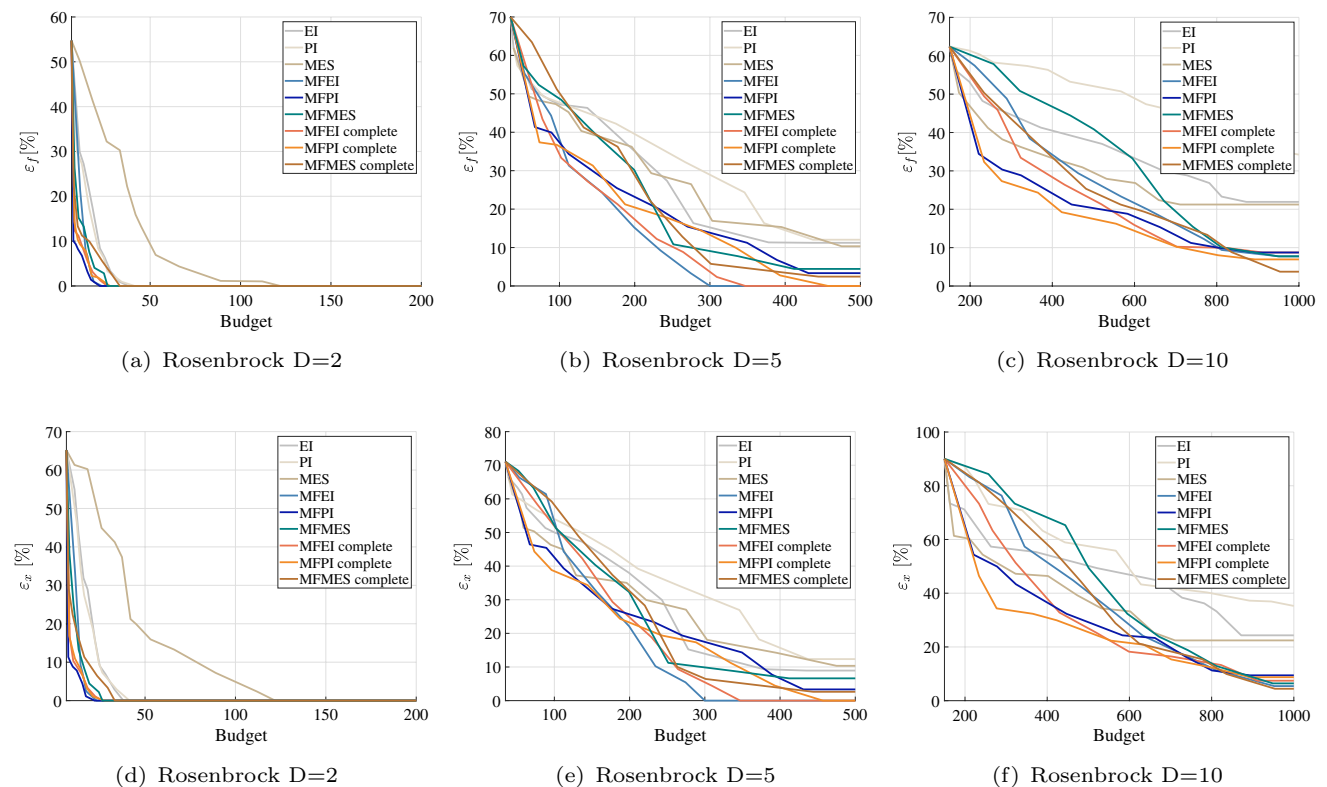


Fig. 12 Performances of the competing algorithms for the Rosenbrock benchmarks

fidelity that pushes the exploration in regions far from the optimum. Second, pure exploitation or a balanced search between exploration and exploitation are advantageous in low-dimensional domains, while pure exploration sacrifices valuable computational resources to improve the awareness about the global distribution of the objective instead of searching the optimum.

Increasing the dimension of the input space to $D = 5$ (Fig. 12b, e), only the MFEI and the MFPI using all the fidelities available are capable to identify the optimum solution, while the other competing algorithms converge to sub-optimal solutions. However, it should be noted the much faster convergence of MFEI with respect to the MFPI in complete setting. These outcomes indicate that as the number of optimization variables increases, both exploration and exploitation are required for an efficient learning procedure. In particular, the exploration improves the accuracy of the surrogate over the domain that better informs the learner during the exploitation phase. The utility of purely exploitation—MFPI—also continues to be observed, but the effectiveness is limited by the dimensionality of the domain that requires an exploration phase to better capture the distribution of the objective function.

Pushing further the dimensionality of the domain at $D = 10$ (Fig. 12c, f), all the algorithms are not capable to reach the analytical optimum with the allocated budget. This can be explained with the unreliable prediction of the surrogate model that is not capable to correctly inform the learner with limited amount of data—limited allocated budget. However, the multifidelity paradigms achieve larger reductions of both the error in the domain ϵ_x and the goal error ϵ_f if compared with the single-fidelity outcomes. This suggests that learners capable to leverage multiple information sources might produce higher gains in a limited budget scenario thanks to the massive use of cheap low-fidelity models to learn the objective function. Among the competing strategies, MFMES exhibits remarkable outcomes in terms of convergence values of the errors when all the library of fidelities is available. This results can be justified with the over-exploration properties of the MFMES acquisition function: the learner uses massive low-fidelity data to refine the approximation of the surrogate model and augment its predictive capabilities. This permits to better inform the procedure and direct computational resources toward the optimum.

The results obtained for the ALOS benchmark problem in Fig. 13 confirm the previous observations about the different effectiveness of the learning schemes. In particular, the multifidelity strategies provides larger accelerations of the optimization procedure in presence oscillations at different frequencies of the objective function for the one- (Fig. 13a, d), two- (Fig. 13b, e) and three- (Fig. 13c, f) dimensional ALOS problem. We observe that the best performances are

delivered by either learners based on the balance between informativeness and representativeness/diversity—MFEI and EI—or a purely informativeness-driven—MFPI and PI—, while over-exploration performs relatively poorly—MFMES and MES. This results are justified with the low-dimensionality of the objective function.

The outcomes related to the multimodal benchmarks are reported in Fig. 14. The multifidelity algorithms are capable to converge toward the analytical optimum with a fraction of the computational cost, if compared with the single-fidelity results. For the Rastrigin function (Fig. 14a, d), the multifidelity methods implementing all the levels of fidelity $l = 1, 2, 3$ outperforms the multifidelity methods with $l = 1, 3$: the intermediate level of fidelity $l = 2$ is more accurate, if compared with the low-fidelity output $l = 3$ and allows to improve the reliability of the Gaussian process in presence of a strong multimodal behaviour. The best performing method is MFPI using $l = 1, 2, 3$ denoting that the over-exploitation of the input space with lower-fidelities levels $l = 2, 3$ allows to take full advantage from low-fidelity data, improving the performance of the learning process. In contrast, we observe that the MES algorithm exhibits a more efficient convergence of the MFMES counterpart. This is related to the already noticed over-exploration of the domain: the MES uses accurate high-fidelity observations to refine the surrogate during the exploration, while the MFMES systematically adopts lower-levels of fidelity to massively query the domain and retard the exploitation with more accurate information sources. The results achieved for the mass spring benchmark problem (Fig. 14b, e) confirm the superior convergence performance of multifidelity algorithms in presence of marked multimodal objective functions. In particular, the balance between exploration and exploitation delivered by the MFEI allows for superior accelerations and contained demand for computational resources. Similar results can be observed for the Paciorek benchmark problem (Fig. 14c, f): the multifidelity learning delivers efficient optimization procedures even in the simultaneous presence of multi-modality and noise. It should be noticed that in presence of noise both the MES and MFMES show an attenuation of the exploratory behaviour and a greater exploitation of the domain. This result is in agreement with what observed by Nguyen et al. [98]. The overall outcomes for these subset of benchmark functions demonstrate that a learning scheme characterized by a balanced exploration and exploitation phases is essential in presence of multimodal behaviour and noise in the measurements of the objective function.

6.3 Advice on using Learning Criteria

Throughout the experiments in this paper and in our research experience, we can summarize several recommendations

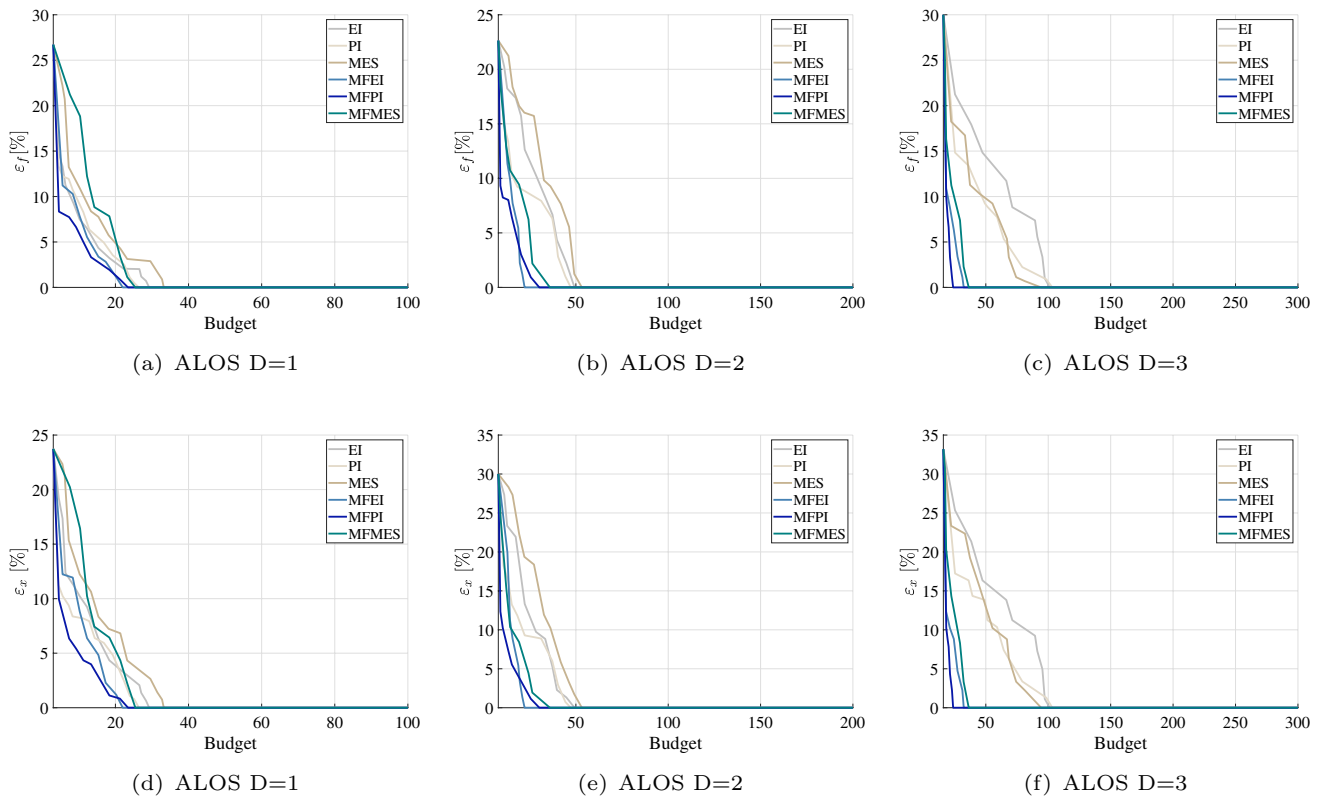


Fig. 13 Performances of the competing algorithms for the ALOS benchmarks

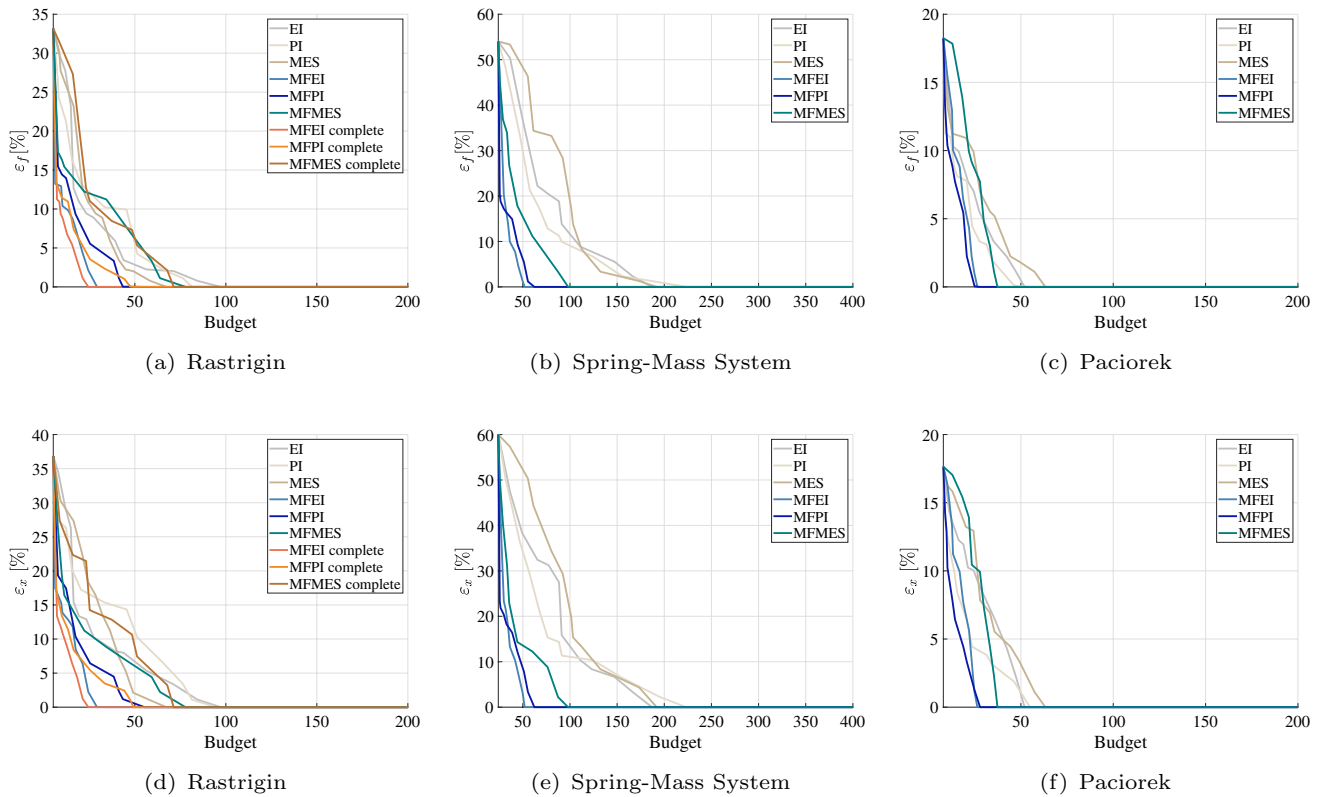


Fig. 14 Performances of the competing algorithms for the multimodal benchmarks

that are intended to provide a guideline to apply the different learning criteria in real-world optimization problems. Although these advice may not be suitable in general due to the vast and natural heterogeneity of the applications where optimization is relevant, we believe that these guidelines can be useful in directing researchers towards the effective use of learning schemes.

1. Pure exploitation/informativeness learning schemes could be potentially beneficial for low-dimensional optimization problems. In our experience, the direct exploitation of data at the beginning of the optimization procedure can produce significant improvement in the solution with relatively contained computational resources. The reason behind this behaviour is due to the accurate prediction of the emulator with contained amount of data in low-dimensional domains. This contributes to better inform the learner and effectively direct resources toward the optimum.
2. Pure exploration/representativeness-diversity could impact considerably the optimization results for high-dimensional optimization problems. The exploration reduces the uncertainty of the emulator over all the domain and leads to a more reliable predictive framework. This would better inform the learner and help directing the computational resources in regions of the domain where is more likely to achieve benefits in terms of solution.
3. The balance between exploration and exploitation guarantees consistent and satisfactory optimization performances over different mathematical properties of the objective function. In particular, our experiments suggest that pursuing the trade-off between exploration and exploitation often leads to satisfactory and in many cases better performance than implementing the learning criteria individually. Although the well performing behaviour in general, it should be privileged mainly in cases when there is no prior knowledge about the specific optimization problem considered to increase the chances of success.
4. When the computational resources are severely limited—e.g. engineering preliminary design phases or trade-off analysis –, there is a clear advantage of using multifidelity learning criteria and leverage a spectrum of information sources at different levels of fidelity. Indeed, the wise combination of fast low-fidelity data with expensive high-fidelity evaluations reduces the overall demand for computational resources, and shows more robust performance for challenging mathematical properties of the objective function such as local/global behaviours, non-linearities and discontinuities, multimodality, and noisy measurements.

7 Concluding Remarks

This paper proposes an original unified perspective of Bayesian optimization and active learning as adaptive sampling schemes guided by common learning principles toward a given optimization goal. Our arguments are based on the recognition of Bayesian optimization and active learning as goal-driven learning procedures characterized by the mutual information exchange between the learner and the surrogate model: the learner makes a decision based on the surrogate information to maximize the sampling utility with respect to the given goal, while the emulator is constantly updated through the results of this decision. Accordingly, we clarify and support our discussion through a general classification of adaptive sampling methodologies, and recognize Bayesian optimization as the logic intersection between active learning and adaptive sampling. This lays the foundations for the explicit formalization of the synergy between Bayesian optimization and active learning considering both a single information source and when a library of representations at different levels of fidelity is available to the learner. This unified perspective is based on the dualism between the active learning criteria of informativeness and representativeness/diversity, and the Bayesian infill criteria of exploration and exploitation as the driving elements to achieve the learning goal. To support our perspective, we reviewed and analysed popular formulations of the acquisition function for Bayesian optimization considering both single-fidelity and multifidelity settings. Accordingly, we formalize this synergy mapping the informativeness learning criterion with the exploitation infill criterion as driving components that direct the selection of samples toward the learning goal. Similarly, we formulate the substantial analogy between representativeness-diversity learning criterion and the exploration infill criterion as sampling policies that improve the awareness about the objective function over the domain. Through stressful analytical benchmark problems, the authors demonstrate the benefits of each learning/infill criteria over challenging mathematical properties of the objective function typically encountered in real-world applications. The results reveal that the balance between the learning/infill criteria ensures good performances and computational efficiency over all the benchmark problems. In addition, multifidelity learning schemes deliver significant accelerations of the learning procedure making them particularly attractive when the available computational resources are limited. The authors also include some advice and guidelines on the use of the different learning criteria based on the experimental results and their own experience in the field.

Acknowledgements This work was supported by project Multi-source Frameworks to Support Real-time Structural Assessment and

Autonomous Decision Making under the Visiting Professor program of Politecnico di Torino, and by the University's Doctoral Scholarship.

Author Contributions FDF and LM contributed to the study conception and design. FDF and MN performed the material preparation, data collection and analysis. The first draft of the manuscript was written by FDF and LM provided supervision, revised, edited and commented on all the versions of the manuscript. All authors approved the final manuscript.

Funding Open access funding provided by Politecnico di Torino within the CRUI-CARE Agreement.

Declarations

Conflict of interest The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

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/>.

References

- Abe N (1998) Query learning strategies using boosting and bagging. In: Proceedings of the 15 international CMF on machine learning (ICML98), pp 1–9
- Atchadé YF, Rosenthal JS (2005) On adaptive Markov chain Monte Carlo algorithms. *Bernoulli* 11(5):815–828
- Atchade Y, Fort G, Moulines E et al (2011) Adaptive Markov chain Monte Carlo: Theory and methods. *Bayesian time series models 1*
- Balakrishnan S, Nguyen QP, Low BKH et al (2020) Efficient exploration of reward functions in inverse reinforcement learning via Bayesian optimization. *Adv Neural Inf Process Syst* 33:4187–4198
- Balcan MF, Broder A, Zhang T (2007) Margin based active learning. In: International conference on computational learning theory. Springer, Berlin, pp 35–50
- Beran PS, Bryson D, Thelen AS et al (2020) Comparison of multi-fidelity approaches for military vehicle design. In: AIAA AVIATION 2020 forum, p 3158
- Berkenkamp F, Krause A, Schoellig AP (2021) Bayesian optimization with safety constraints: safe and automatic parameter tuning in robotics. *Mach Learn* 112:1–35
- Bonfiglio L, Perdikaris P, Brizzolaro S et al (2018) Multi-fidelity optimization of super-cavitating hydrofoils. *Comput Methods Appl Mech Eng* 332:63–85
- Bugallo MF, Martino L, Corander J (2015) Adaptive importance sampling in signal processing. *Digi Signal Process* 47:36–49
- Bugallo MF, Elvira V, Martino L et al (2017) Adaptive importance sampling: the past, the present, and the future. *IEEE Signal Process Mag* 34(4):60–79
- Bui-Thanh T, Willcox K, Ghattas O et al (2007) Goal-oriented, model-constrained optimization for reduction of large-scale systems. *J Comput Phys* 224(2):880–896
- Burbidge R, Rowland JJ, King RD (2007) Active learning for regression based on query by committee. In: International conference on intelligent data engineering and automated learning. Springer, Berlin, pp 209–218
- Cai W, Zhang Y, Zhou J (2013) Maximizing expected model change for active learning in regression. In: 2013 IEEE 13th international conference on data mining, IEEE, pp 51–60
- Cai W, Zhang M, Zhang Y (2016) Batch mode active learning for regression with expected model change. *IEEE Trans Neural Netw Learn Syst* 28(7):1668–1681
- Chattopadhyay R, Wang Z, Fan W et al (2013) Batch mode active sampling based on marginal probability distribution matching. *ACM Trans Knowl Discov Data (TKDD)* 7(3):1–25
- Chernoff H (1959) Sequential design of experiments. *Ann Math Stat* 30(3):755–770
- Cohn D (1993) Neural network exploration using optimal experiment design. *Adv Neural Inf Process Syst* 6
- Dasgupta S, Hsu D (2008) Hierarchical sampling for active learning. In: Proceedings of the 25th international conference on Machine learning, pp 208–215
- Davis SE, Cremaschi S, Eden MR (2018) Efficient surrogate model development: impact of sample size and underlying model dimensions. In: Computer aided chemical engineering, vol 44. Elsevier, New York, pp 979–984
- Demir B, Bruzzone L (2014) A multiple criteria active learning method for support vector regression. *Pattern Recognit* 47(7):2558–2567
- Di Fiore F, Mainini L (2022) Non-myopic multifidelity Bayesian optimization. [arXiv:2207.06325](https://arxiv.org/abs/2207.06325)
- Di Fiore F, Mainini L (2023) Nm-mf: Non-myopic multifidelity framework for constrained multi-regime aerodynamic optimization. *AIAA J* 61(3):1270–1280
- Di Fiore F, Maggiore P, Mainini L (2021) Multifidelity do-main-aware learning for the design of re-entry vehicles. *Struct Multi-discip Optim* 64(5):3017–3035
- Dias L, Bhosekar A, Ierapetritou M (2019) Adaptive sampling approaches for surrogate-based optimization. *Computer aided chemical engineering*, vol 47. Elsevier, New York, pp 377–384
- Donmez P, Carbonell JG, Schneider J (2009) Efficiently learning the accuracy of labeling sources for selective sampling. In: Proceedings of the 15th ACM SIGKDD international conference on Knowledge discovery and data mining, pp 259–268
- Dudley JJ, Jacques JT, Kristensson PO (2019) Crowdsourcing interface feature design with Bayesian optimization. In: Proceedings of the 2019 CHI conference on human factors in computing systems, pp 1–12
- Eigel M, Ernst OG, Sprungk B et al (2022) On the convergence of adaptive stochastic collocation for elliptic partial differential equations with affine diffusion. *SIAM J Numer Anal* 60(2):659–687
- Elder JF (1992) Global $r/\sup d$ /optimization when probes are expensive: the grope algorithm. In: [Proceedings] 1992 IEEE international conference on systems, man, and cybernetics, IEEE, pp 577–582
- El-Laham Y, Martino L, Elvira V et al (2019) Efficient adaptive multiple importance sampling. In: 2019 27th European signal processing conference (EUSIPCO), IEEE, pp 1–5
- Eriksson D, Pearce M, Gardner J et al (2019) Scalable global optimization via local Bayesian optimization. *Adv Neural Inf Process Syst* 32
- Fang M, Yin J, Tao D (2014) Active learning for crowdsourcing using knowledge transfer. In: Proceedings of the AAAI conference on artificial intelligence

32. Forrester AI, Sóbester A, Keane AJ (2007) Multi-fidelity optimization via surrogate modelling. *Proc R Soc A* 463(2088):3251–3269
33. Frazier PI (2018) A tutorial on Bayesian optimization. [arXiv:1807.02811](https://arxiv.org/abs/1807.02811)
34. Frazier PI, Wang J (2016) Bayesian optimization for materials design. *Inf Sci Mater Discov Des* 225:45–75
35. Fu Y, Li B, Zhu X et al (2013) Active learning without knowing individual instance labels: a pairwise label homogeneity query approach. *IEEE Trans Knowl Data Eng* 26(4):808–822
36. Gao R, Saar-Tsechansky M (2020) Cost-accuracy aware adaptive labeling for active learning. In: *Proceedings of the AAAI conference on artificial intelligence*, pp 2569–2576
37. Gerstner T, Griebel M (2003) Dimension-adaptive tensor-product quadrature. *Computing* 71:65–87
38. Giselle Fernández-Godino M, Park C, Kim NH et al (2019) Issues in deciding whether to use multifidelity surrogates. *AIAA J* 57(5):2039–2054
39. Gonzalez J, Lezmi E, Roncalli T et al (2019) Financial applications of gaussian processes and Bayesian optimization. [arXiv:1903.04841](https://arxiv.org/abs/1903.04841)
40. Grassi F, Manganini G, Garraffa M et al (2023) RAAL: resource aware active learning for multifidelity efficient optimization. *AIAA J* 61(6):2744–2753
41. Guignard D, Nobile F (2018) A posteriori error estimation for the stochastic collocation finite element method. *SIAM J Numer Anal* 56(5):3121–3143
42. Gunst RF, Mason RL (2009) *Fractional factorial design*. Wiley Interdiscip Rev 1(2):234–244
43. Gunzburger M, Webster CG, Zhang G (2014) An adaptive wavelet stochastic collocation method for irregular solutions of partial differential equations with random input data. In: *Sparse grids and applications-Munich 2012*, Springer, pp 137–170
44. Gustafsson A, Herrmann A, Huber F (2013) *Conjoint measurement: methods and applications*. Springer, New York
45. Haji-Ali AL, Nobile F, Tamellini L et al (2016) Multi-index stochastic collocation for random PDES. *Comput Methods Appl Mech Eng* 306:95–122
46. He T, Zhang S, Xin J et al (2014) An active learning approach with uncertainty, representativeness, and diversity. *Sci World J*. <https://doi.org/10.1155/2014/827586>
47. Hennig P, Schuler CJ (2012) Entropy search for information-efficient global optimization. *J Mach Learn Res* 13:1809–1837
48. Hesterberg T (1995) Weighted average importance sampling and defensive mixture distributions. *Technometrics* 37(2):185–194
49. Holub A, Perona P, Burl MC (2008) Entropy-based active learning for object recognition. In: *2008 IEEE computer society conference on computer vision and pattern recognition workshops*, IEEE, pp 1–8
50. Hsu WN, Lin HT (2015) Active learning by learning. In: *Twenty-ninth AAAI conference on artificial intelligence*
51. Huang D, Allen TT, Notz WI et al (2006) Sequential kriging optimization using multiple-fidelity evaluations. *Struct Multidiscip Optim* 32(5):369–382
52. Huang SJ, Chen JL, Mu X et al (2017) Cost-effective active learning from diverse labelers. In: *IJCAI*, pp 1879–1885
53. Imani M, Ghoreishi SF (2021) Scalable inverse reinforcement learning through multifidelity Bayesian optimization. *IEEE Trans Neural Netw Learn Syst* 33(8):4125–4132
54. Ipeirotis PG, Provost F, Sheng VS et al (2014) Repeated labeling using multiple noisy labelers. *Data Min Knowl Discov* 28(2):402–441
55. Irshad F, Karsch S, Döpp A (2023) Multi-objective and multifidelity Bayesian optimization of laser-plasma acceleration. *Phys Rev Res* 5(1):013063
56. Jakeman JD, Roberts SG (2012) Local and dimension adaptive stochastic collocation for uncertainty quantification. In: *Sparse grids and applications*. Springer, Berlin, pp 181–203
57. Jakeman JD, Eldred MS, Geraci G et al (2020) Adaptive multi-index collocation for uncertainty quantification and sensitivity analysis. *Int J Numer Methods Eng* 121(6):1314–1343
58. Jensen H, Jerez D, Valdebenito M (2020) An adaptive scheme for reliability-based global design optimization: a Markov chain Monte Carlo approach. *Mech Syst Signal Process* 143:106836
59. Jin R, Chen W, Sudjianto A (2002) On sequential sampling for global metamodeling in engineering design. In: *International design engineering technical conferences and computers and information in engineering conference*, pp 539–548
60. Jones DR (2001) A taxonomy of global optimization methods based on response surfaces. *J Global Optim* 21(4):345–383
61. Jones DR, Schonlau M, Welch WJ (1998) Efficient global optimization of expensive black-box functions. *J Global Optim* 13(4):455–492
62. Kandasamy K, Schneider J, Póczos B (2015) High dimensional Bayesian optimisation and bandits via additive models. In: *International conference on machine learning*, PMLR, pp 295–304
63. Kandasamy K, Dasarathy G, Schneider J et al (2017) Multifidelity Bayesian optimisation with continuous approximations. In: *International conference on machine learning*, PMLR, pp 1799–1808
64. Karamchandani A, Bjerager P, Cornell C (1989) Adaptive importance sampling. In: *Structural safety and reliability*. ASCE, pp 855–862
65. Kennedy MC, O’Hagan A (2000) Predicting the output from a complex computer code when fast approximations are available. *Biometrika* 87(1):1–13
66. Kim NH, An D, Choi JH (2017) *Prognostics and health management of engineering systems*. Springer, Switzerland
67. Kong H, Yan J, Wang H et al (2020) Energy management strategy for electric vehicles based on deep q-learning using Bayesian optimization. *Neural Comput Appl* 32:14431–14445
68. Koyama Y, Sato I, Goto M (2020) Sequential gallery for interactive visual design optimization. *ACM Trans Graph* 39(4):88–91
69. Kushner HJ (1964) A new method of locating the maximum point of an arbitrary multiplex curve in the presence of noise. *ASME J Basic Eng* 86:97–106
70. Lam R, Willcox K (2017) Lookahead Bayesian optimization with inequality constraints. *Adv Neural Inf Process Syst* 30
71. Lam R, Poloczek M, Frazier P et al (2018) Advances in Bayesian optimization with applications in aerospace engineering. In: *2018 AIAA non-deterministic approaches conference*, p 1656
72. Lang J, Scheichl R, Silvester D (2020) A fully adaptive multilevel stochastic collocation strategy for solving elliptic PDEs with random data. *J Comput Phys* 419:109692
73. Lazin MF, Shelton CR, Sandhofer SN et al (2023) High-dimensional multi-fidelity Bayesian optimization for quantum control. *Mach Learn* 4(4):045014
74. Lewis DD, Catlett J (1994) Heterogeneous uncertainty sampling for supervised learning. In: *Machine learning proceedings 1994*. Elsevier, New York, pp 148–156
75. Li J, Aickelin U (2003) A Bayesian optimization algorithm for the nurse scheduling problem. In: *The 2003 congress on evolutionary computation, 2003. CEC’03.*, IEEE, pp 2149–2156
76. Li X, Guo Y (2013) Adaptive active learning for image classification. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp 859–866
77. Li M, Sethi IK (2006) Confidence-based active learning. *IEEE Trans Pattern Anal Mach Intell* 28(8):1251–1261
78. Lieberman C, Willcox K (2013) Goal-oriented inference: approach, linear theory, and application to advection diffusion. *SIAM Rev* 55(3):493–519

79. Liu H, Ong YS, Cai J (2018) A survey of adaptive sampling for global metamodeling in support of simulation-based complex engineering design. *Struct Multidiscip Optim* 57:393–416
80. Liu Y, Chen S, Wang F et al (2018) Sequential optimization using multi-level cokriging and extended expected improvement criterion. *Struct Multidiscip Optim* 58(3):1155–1173
81. Lu L, Meng X, Mao Z et al (2021) Deepxde: a deep learning library for solving differential equations. *SIAM Rev* 63(1):208–228
82. Ma X, Zabaras N (2009) An efficient Bayesian inference approach to inverse problems based on an adaptive sparse grid collocation method. *Inverse Prob* 25(3):035013
83. Mainini L, Serani A, Rumpfkeil M et al (2022) Analytical benchmark problems for multifidelity optimization methods. [arXiv:2204.07867](https://arxiv.org/abs/2204.07867)
84. Martino L, Elvira V, Luengo D et al (2015) An adaptive population importance sampler: learning from uncertainty. *IEEE Trans Signal Process* 63(16):4422–4437
85. Martins JR, Ning A (2021) *Engineering design optimization*. Cambridge University Press, Cambridge
86. Marzouk Y, Xiu D (2009) A stochastic collocation approach to Bayesian inference in inverse problems. *Commun Comput Phys* 6(4):826–847
87. McKay MD (1992) Latin hypercube sampling as a tool in uncertainty analysis of computer models. In: *Proceedings of the 24th conference on Winter simulation*, pp 557–564
88. McKay MD, Beckman RJ, Conover WJ (2000) A comparison of three methods for selecting values of input variables in the analysis of output from a computer code. *Technometrics* 42(1):55–61
89. Meliani M, Bartoli N, Lefebvre T et al (2019) Multi-fidelity efficient global optimization: methodology and application to airfoil shape design. In: *AIAA aviation 2019 forum*, p 3236
90. Moćkus J (1974) On Bayesian methods for seeking the extremum. In: *Optimization techniques IFIP technical conference: novosibirsk, July 1–7*. Springer, Berlin, pp 400–404
91. Mockus J (1975) On the Bayes methods for seeking the extremal point. *IFAC Proc Vol* 8(1):428–431
92. Mockus J (2012) *Bayesian approach to global optimization: theory and applications*, vol 37. Springer, Berlin
93. Monarch RM (2021) *Human-in-the-loop machine learning: active learning and annotation for human-centered AI*. Simon and Schuster, New York
94. Montgomery DC (2017) *Design and analysis of experiments*. Wiley, New York
95. Myers RH, Montgomery DC, Anderson-Cook CM (2016) *Response surface methodology: process and product optimization using designed experiments*. Wiley, New York
96. Nabian MA, Gladstone RJ, Meidani H (2021) Efficient training of physics-informed neural networks via importance sampling. *Comput Aided Civil Infrastruct Eng* 36(8):962–977
97. Nayebi A, Munteanu A, Poloczek M (2019) A framework for Bayesian optimization in embedded subspaces. In: *International conference on machine learning*, PMLR, pp 4752–4761
98. Nguyen QP, Low BKH, Jaillet P (2022) Rectified max-value entropy search for Bayesian optimization. [arXiv:2202.13597](https://arxiv.org/abs/2202.13597)
99. Oden JT, Vemaganti KS (2000) Estimation of local modeling error and goal-oriented adaptive modeling of heterogeneous materials: I. Error estimates and adaptive algorithms. *J Comput Phys* 164(1):22–47
100. Osborne MA, Garnett R, Roberts SJ (2009) Gaussian processes for global optimization. In: *3rd international conference on learning and intelligent optimization (LION3)*, Citeseer, pp 1–15
101. Owen AB (2003) Quasi-Monte Carlo sampling. *Monte Carlo Ray Tracing* 1:69–88
102. Peherstorfer B, Cui T, Marzouk Y et al (2016) Multifidelity importance sampling. *Comput Methods Appl Mech Eng* 300:490–509
103. Peherstorfer B, Willcox K, Gunzburger M (2018) Survey of multifidelity methods in uncertainty propagation, inference, and optimization. *SIAM Rev* 60(3):550–591
104. Perdikaris P, Karniadakis GE (2016) Model inversion via multi-fidelity Bayesian optimization: a new paradigm for parameter estimation in haemodynamics, and beyond. *J R Soc Interface* 13(118):20151107
105. Pezzuto S, Perdikaris P, Costabal FS (2022) Learning cardiac activation maps from 12-lead eeg with multi-fidelity Bayesian optimization on manifolds. [arXiv:2203.06222](https://arxiv.org/abs/2203.06222)
106. Pour ES, Jafari H, Lashgari A et al (2022) Cryptocurrency price prediction with neural networks of LSTM and Bayesian optimization. *Eur J Bus Manag Res* 7(2):20–27
107. Priem R, Gagnon H, Chittick I et al (2020) An efficient application of Bayesian optimization to an industrial MDO framework for aircraft design. In: *AIAA aviation 2020 forum*, p 3152
108. Provost F, Jensen D, Oates T (1999) Efficient progressive sampling. In: *Proceedings of the fifth ACM SIGKDD international conference on Knowledge discovery and data mining*, pp 23–32
109. Ram A, Leake DB (1995) *Goal-driven learning*. MIT Press, Cambridge
110. Rasmussen CE (2003) *Gaussian processes in machine learning*. Summer school on machine learning. Springer, New York, pp 63–71
111. RayChaudhuri T, Hamey LG (1995) Minimisation of data collection by active learning. In: *Proceedings of ICNN'95-international conference on neural networks*, IEEE, pp 1338–1341
112. Reitmaier T, Calma A, Sick B (2015) Transductive active learning—a new semi-supervised learning approach based on iteratively refined generative models to capture structure in data. *Inf Sci* 293:275–298
113. Robert CP, Casella G, Casella G (1999) *Monte Carlo statistical methods*, vol 2. Springer, Berlin
114. Ruan X, Jiang P, Zhou Q et al (2020) Variable-fidelity probability of improvement method for efficient global optimization of expensive black-box problems. *Struct Multidiscip Optim* 62(6):3021–3052
115. Sarkar S, Mondal S, Joly M et al (2019) Multifidelity and multiscale Bayesian framework for high-dimensional engineering design and calibration. *J Mech Des* 141(12):121001
116. Scott W, Frazier P, Powell W (2011) The correlated knowledge gradient for simulation optimization of continuous parameters using gaussian process regression. *SIAM J Optim* 21(3):996–1026
117. Sener O, Savarese S (2017) A geometric approach to active learning for convolutional neural networks. [arXiv:1708.00489](https://arxiv.org/abs/1708.00489)
118. Serani A, Leotardi C, Iemma U et al (2016) Parameter selection in synchronous and asynchronous deterministic particle swarm optimization for ship hydrodynamics problems. *Appl Soft Comput* 49:313–334
119. Serani A, Ficini S, Grigoropoulos G et al (2022) Resistance and seakeeping optimization of a naval destroyer by multi-fidelity methods. *VCG* 1000(1.390):0–056
120. Settles B (2009) *Active learning literature survey*. Technical Report TR-1648
121. Settles B, Craven M (2008) An analysis of active learning strategies for sequence labeling tasks. In: *proceedings of the 2008 conference on empirical methods in natural language processing*, pp 1070–1079
122. Seung HS, Opper M, Sompolinsky H (1992) Query by committee. In: *Proceedings of the fifth annual workshop on computational learning theory*, pp 287–294

123. Shah A, Ghahramani Z (2015) Parallel predictive entropy search for batch global optimization of expensive objective functions. *Adv Neural Inf Process Syst* 28
124. Shahriari B, Swersky K, Wang Z et al (2015) Taking the human out of the loop: a review of Bayesian optimization. *Proc IEEE* 104(1):148–175
125. Shapiro A (2003) Monte Carlo sampling methods. *Handb Oper Res Manag Sci* 10:353–425
126. Shen D, Zhang J, Su J et al (2004) Multi-criteria-based active learning for named entity recognition. In: *Proceedings of the 42nd annual meeting of the Association for Computational Linguistics (ACL-04)*, pp 589–596
127. Sobester A, Forrester A, Keane A (2008) *Engineering design via surrogate modelling: a practical guide*. Wiley, New York
128. Stuckman BE (1988) A global search method for optimizing non-linear systems. *IEEE Trans Syst Man Cybern* 18(6):965–977
129. Sugiyama M, Nakajima S (2009) Pool-based active learning in approximate linear regression. *Mach Learn* 75(3):249–274
130. Takeno S, Fukuoka H, Tsukada Y et al (2020) Multi-fidelity Bayesian optimization with max-value entropy search and its parallelization. In: *International conference on machine learning*, PMLR, pp 9334–9345
131. Tang YP, Huang SJ (2019) Self-paced active learning: query the right thing at the right time. In: *Proceedings of the AAAI conference on artificial intelligence*, pp 5117–5124
132. Turner R, Eriksson D, McCourt M et al (2021) Bayesian optimization is superior to random search for machine learning hyperparameter tuning: analysis of the black-box optimization challenge 2020. In: *NeurIPS 2020 competition and demonstration track*, PMLR, pp 3–26
133. Viana FA, Simpson TW, Balabanov V et al (2014) Special section on multidisciplinary design optimization: metamodeling in multidisciplinary design optimization: how far have we really come? *AIAA J* 52(4):670–690
134. Victoria AH, Maragatham G (2021) Automatic tuning of hyperparameters using Bayesian optimization. *Evol Syst* 12:217–223
135. Wang Z, Jegelka S (2017) Max-value entropy search for efficient Bayesian optimization. In: *International conference on machine learning*, PMLR, pp 3627–3635
136. Wang Z, Ye J (2015) Querying discriminative and representative samples for batch mode active learning. *ACM Trans Knowl Discov Data* 9(3):1–23
137. Wang Z, Hutter F, Zoghi M et al (2016) Bayesian optimization in a billion dimensions via random embeddings. *J Artif Intell Res* 55:361–387
138. Wang Z, Gehring C, Kohli P et al (2018) Batched large-scale Bayesian optimization in high-dimensional spaces. In: *International conference on artificial intelligence and statistics*, PMLR, pp 745–754
139. Wang J, Clark SC, Liu E et al (2020) Parallel Bayesian global optimization of expensive functions. *Oper Res* 68(6):1850–1865
140. Winter J, Abaidi R, Kaiser J et al (2023) Multi-fidelity Bayesian optimization to solve the inverse Stefan problem. *Comput Methods Appl Mech Eng* 410:115946
141. Wu D (2018) Pool-based sequential active learning for regression. *IEEE Trans Neural Netw Learn Syst* 30(5):1348–1359
142. Wu J, Frazier P (2019) Practical two-step lookahead Bayesian optimization. *Adv Neural Inf Process Syst* 32
143. Wu TP, Tang CK (2005) A Bayesian approach for shadow extraction from a single image. In: *Tenth IEEE international conference on computer vision (ICCV'05)*, vol 1, pp 480–487
144. Wu D, Lawhern VJ, Gordon S et al (2016) Offline eeg-based driver drowsiness estimation using enhanced batch-mode active learning (EBMAL) for regression. In: *2016 IEEE international conference on systems, man, and cybernetics (SMC)*, IEEE, pp 000730–000736
145. Wu J, Toscano-Palmerin S, Frazier PI et al (2020) Practical multi-fidelity Bayesian optimization for hyperparameter tuning. In: *Uncertainty in artificial intelligence*, PMLR, pp 788–798
146. Wu C, Zhu M, Tan Q et al (2023) A comprehensive study of non-adaptive and residual-based adaptive sampling for physics-informed neural networks. *Comput Methods Appl Mech Eng* 403:115671
147. Xiao NC, Zhan H, Yuan K (2020) A new reliability method for small failure probability problems by combining the adaptive importance sampling and surrogate models. *Comput Methods Appl Mech Eng* 372:113336
148. Yan Y, Rosales R, Fung G et al (2011) Active learning from crowds. In: *ICML*
149. Yan Y, Rosales R, Fung G et al (2012) Active learning from multiple knowledge sources. In: *Artificial intelligence and statistics*, PMLR, pp 1350–1357
150. Young MT, Hinkle JD, Kannan R et al (2020) Distributed Bayesian optimization of deep reinforcement learning algorithms. *J Parallel Distrib Comput* 139:43–52
151. Yu G, Chen X, Domeniconi C et al (2020) CMAL: cost-effective multi-label active learning by querying subexamples. *IEEE Trans Knowl Data Eng* 34(5):2091–2105
152. Zang TA (2002) *Needs and opportunities for uncertainty-based multidisciplinary design methods for aerospace vehicles*. National Aeronautics and Space Administration, Langley Research Center
153. Zhan X, Liu H, Li Q et al (2021) A comparative survey: Benchmarking for pool-based active learning. In: *IJCAI*, pp 4679–4686
154. Zhang Y, Hoang TN, Low BKH et al (2017) Information-based multi-fidelity Bayesian optimization. In: *NIPS workshop on Bayesian optimization*
155. Zhao Y, Xu C, Cao Y (2006) Research on query-by-committee method of active learning and application. In: *International conference on advanced data mining and applications*, Springer, pp 985–991
156. Zhao L, Sukthankar G, Sukthankar R (2011) Incremental relabeling for active learning with noisy crowdsourced annotations. In: *2011 IEEE third international conference on privacy, security, risk and trust and 2011 IEEE third international conference on social computing*, IEEE, pp 728–733
157. Zhao Y, Shi Z, Zhang J et al (2019) A novel active learning framework for classification: Using weighted rank aggregation to achieve multiple query criteria. *Pattern Recognit* 93:581–602
158. Zhilinskis A (1975) Single-step Bayesian search method for an extremum of functions of a single variable. *Cybernetics* 11(1):160–166
159. Zhou X, Lu Y, Lu J et al (2011) Abrupt motion tracking via intensively adaptive Markov-chain Monte Carlo sampling. *IEEE Trans Image Process* 21(2):789–801

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