

Adaptive Simulation-Based Training of Artificial-Intelligence Decision Makers Using Bayesian Optimization

Brett Israelsen* and Nisar Ahmed†

University of Colorado, Boulder, Colorado 80309

Kenneth Center‡ and Roderick Green§

Orbit Logic Incorporated, Greenbelt, Maryland 20770

and

Winston Bennett, Jr.¶

U.S. Air Force Research Laboratory, Wright Patterson Air Force Base, Ohio 45433

DOI: 10.2514/1.I010553

This work studies how an artificial-intelligence-controlled dogfighting agent with tunable decision-making parameters can learn to optimize performance against an intelligent adversary, as measured by a stochastic objective function evaluated on simulated combat engagements. Gaussian process Bayesian optimization techniques are developed to automatically learn global Gaussian process surrogate models, which provide statistical performance predictions in both explored and unexplored areas of the parameter space. This allows a learning engine to sample full-combat simulations at parameter values that are most likely to optimize performance and provide highly informative data points for improving future predictions. However, standard Gaussian process Bayesian optimization methods do not provide a reliable surrogate model for the highly volatile objective functions found in aerial combat and thus do not reliably identify global maxima. These issues are addressed by novel repeat sampling and hybrid repeat/multipoint sampling techniques. Simulation studies show that hybrid repeat/multipoint sampling improves the accuracy of Gaussian process surrogate models, allowing artificial-intelligence decision makers to more accurately predict performance and efficiently tune parameters.

I. Introduction

Rapid advancement in the capabilities of artificial intelligence (AI) has the potential to completely revolutionize the way that U.S. armed forces train for battlespace dominance. As AI becomes more sophisticated, there are many opportunities to insert AI into domain training environments. One use of AI is as the basis of agents that stand in for human vehicle/platform controllers in live–virtual–constructive (LVC) simulations. Humans participating in training exercises in these environments need to be challenged at an appropriate level relative to their skills. Needed are agents that can assess the skill level of human participants and adapt accordingly, serving as credible and adaptable adversaries that are indistinguishable from experienced humans. This work investigates how an AI agent with tunable parameters governing its overall behavior can be adapted to optimize an objective function for engagement outcomes.

Several challenges make it difficult to meet these objectives.

1) Simulating an engagement can be costly. Beyond the financial expense of operating the simulation environment, contributions to the cost may also include the involvement of skilled personnel with limited availability and the wall-clock duration of the simulation itself.

2) The engagement metrics that need to be optimized cannot be described analytically but can only be evaluated by running simulations. When performance evaluations are sampled, they are generally highly nonlinear functions of environmental parameters and decision-maker states. Consequently, many traditional optimization methods are not applicable.

3) The realistic nature of engagement simulations makes virtually all performance objective functions of interest extremely volatile and uncertain (e.g., due to combined random effects of weather, terrain, sensor noise, psychomotor time delays, etc.).

Gaussian process Bayesian optimization (GPBO) surrogate models are well suited for directly addressing points 1 and 2. In this formulation, the Gaussian process (GP) serves as a tractable surrogate model that approximates the true (non-closed-form/intractable) objective function. This surrogate model estimates a nonparametric probability distribution over the objective function values at all location in the solution space. The Bayesian optimization algorithm uses this surrogate model to intelligently search the solution space for the optimum, based on a number of sampled function evaluations (i.e., using “explore–exploit” strategies to locate the minimum as quickly as possible while also using sparse function evaluations to build the surrogate model). Because the GP surrogate model is cheaper to evaluate than the true objective function, global nonlinear optimization methods can be used on the GP model to efficiently search the decision parameter space while also accounting for uncertainty in the underlying objective function.

However, standard GPBO methods are not well suited to address point 3. This work develops a novel approach for implementing GPBO, called hybrid repeat/multipoint sampling (HRMS), to address these issues. In the setting of simulated one-on-one aerial dogfighting engagements, GPBO with HRMS is able to not only identify the optimum more reliably than standard GPBO but also yield a more accurate and consistent surrogate representation of the objective surface, using no more total function evaluations than traditional GPBO techniques.

The remainder of this paper is outlined as follows. Section II discusses relevant prior work and provides a formal definition of the adaptive AI agent problem; relevant details regarding the application of GPBO for aerial dogfighting performance optimization are also provided. Details of

Presented as Paper 2017-0343 at the AIAA Information Systems-AIAA Infotech @ Aerospace, AIAA SciTech Forum, Grapevine, TX, 8–12 January 2017; received 26 March 2017; revision received 29 July 2017; accepted for publication 15 December 2017; published online 30 January 2018. Copyright © 2018 by the American Institute of Aeronautics and Astronautics, Inc. All rights reserved. All requests for copying and permission to reprint should be submitted to CCC at www.copyright.com; employ the ISSN 2327-3097 (online) to initiate your request. See also AIAA Rights and Permissions www.aiaa.org/randp.

*Graduate Researcher, Computer Science. Student Member AIAA.

†Assistant Professor, Aerospace Engineering Sciences. Member AIAA.

‡Director. Professional Member AIAA.

§Senior Software Engineer.

¶Division Technical Advisor, Warfighter Readiness Research Division.

the proposed GPBO “learning engine” framework used to train the decision-making AI are given in Sec. III, which also includes some discussion regarding practical implementation of GPBO. Section IV reviews existing sampling strategies for GPBO and introduces our novel sampling strategy: HRMS. Finally, Sec. V shows by simulated experiments that the proposed implementation of GPBO with HRMS is useful for optimizing highly volatile performance metrics for the dogfighting application. It is shown that HRMS can significantly outperform traditional GPBO sampling techniques when dealing with highly volatile objective functions and yields valuable insights about AI decision-maker performance through the global GP surrogate model.

II. Preliminaries and Problem Description

A. Problem Domain and Previous Work

Logistical and fiscal constraints have led to a recent surge in interest around simulation-based methods for training warfighters. As such, efforts such as the U.S. Air Force’s “Not-So-Grand Challenge” were developed with the specific goals of investigating solutions for current and future simulation training systems. As part of this effort, different autonomous decision-making AI agents have been developed and evaluated based on their ability to effectively mimic human pilots in different situations [1]. Although such AIs can effectively mimic human behavior to varying degrees of success in different circumstances, there still remains the question of whether they can adapt based on their adversaries’ responses. The specific problem considered here is how a single autonomous agent with certain behavioral parameters can adapt to improve performance in response to other autonomous agents (human or AI) in a combat simulation featuring stochastic uncertainties and highly volatile outcomes.

The application focus is on simulations for air-to-air combat (dogfighting) training, which have been studied extensively. For instance, McManus and Goodrich [2] discuss the integration of an AI-based tactical decision generator (TDG) system into two separate simulators to study and evaluate air combat environments. One of the simulation modes included an interface for human pilots to participate in training against the TDG. More recently, the Not-So-Grand Challenge produced multiple “humanlike” AI systems to train against human pilots [1]. Several different teams participated and were each rated on several metrics of performance and effectiveness against human pilots. State-of-the-art techniques, such as inverse reinforcement learning, hierarchical logic, and other proprietary approaches, were used to develop the AI pilots. One crucial factor for performance evaluation in this case is the degree to which the AI can mimic realistic human decision making. Although it was found that some AI systems were more effective than others in this regard for specific scenarios (e.g., execution of evasive maneuvers, formation maintenance, target engagement, etc.), the question of how to efficiently “retune” and adapt any particular AI based on adversarial responses remained open.

This leads to the consideration of another critical component in the evaluation process for pilots (human or AI): development of suitable metrics that quantify performance during an engagement. For instance, Moore et al. discuss formal methods for measuring pilot dogfighting performance and validated it in simulated combat scenarios [3]. These measures of pilot performance are meant to be less subjective than traditional ratings given by instructors/expert observers (as in the Not-So-Grand Challenge) because they are based on data recorded during an engagement. This touches on another motivation for developing AI pilots (i.e., automation of instructor functions and expert training resources, which are costly to provide and maintain). Ideally, if an autonomous AI can control a complex optionally manned vehicle such as a fighter plane, then it should also be able to evaluate and advise human pilots based on that same expertise (in much the same way human instructors are able to do this).

Performance metrics of aircraft engagement scenarios have evolved considerably since the inception of engagement debriefings. Many of these metrics are well accepted in the community. In contemporary development of newer metrics, expert evaluations are still used for validation. Kelly reviews and summarizes much of this work [4] and specifically mentions metrics that include variables such as relative aircraft position, throttle and speedbrake manipulation, and overall engagement outcomes, to name a few. Identification of meaningful metrics that operate on time-series and summary data from engagements is still an active field of research [2,4–8]. The work reported in this paper uses several metrics developed in the fighter-pilot training literature as objective functions for automated learning and tuning of decision-making AI (although it is not exclusive to any particular set of metrics or dogfighting training applications, *per se*).

A large segment of work on optimization of aircraft engagement focuses on optimal teaming strategies. Mulgund et al. examined “large-scale” air combat tactics (formations, etc.) and were able to demonstrate promising results in that area [5,7]. Wu et al. addressed the problem of optimizing cooperative multiple target attack using genetic algorithms (GAs) [9]. Also applying GAs, Gonsalves and Burge investigated how mission plans could be optimized [10]. Although these are interesting and important application areas, the present work is focused on one-on-one engagements between autonomous adversaries (where one or both agents are an AI decision maker) and providing an adversary that can be adaptive to the skills of the other pilot (human/AI). In other words, as opposed to modifying team tactics/strategies, this work focuses on optimally adapting the behaviors of individual agents. Another key difference from this work is that direct optimization strategies such as GAs are not well suited to simulations that feature noisy environments and highly volatile outcomes. The approach proposed in this work overcomes this limitation by producing models of the objective function that account for these sources of uncertainty to produce more reliable performance optimization results.

More recently, Ernest et. al designed an AI system that can function in real time during aerial combat [11]. Their work used genetic optimization of a group of fuzzy inference systems (FISs) that together produce control actions given input data from the simulated environment. This system is able to produce “fine-grained” control actions of a single jet or groups of fighters. Also, because of the specification of the FISs, the control outputs are interpretable to humans. As in the present paper, the work of [11] seeks an optimal solution to a combat scenario by modifying fighter behaviors. However, a key difference is that their approach requires designing the entire AI system from the top down, whereas the approach advocated in this paper operates on top of any AI with tunable behavior parameters. Furthermore, the approach presented here tries to minimize the number of simulations needed to optimally adapt the AI’s behavior to maximize performance, while building an easily interpretable global statistical model of the AI parameter-to-performance metric mapping at the same time. Autonomous AI agent technology has also been extensively developed for the video game industry, where nonplayer characters interact with human players. Still, the majority of agent AIs used in the gaming industry are highly scripted and nonadaptive [12]. Cole et al. used GAs to tune game agents for first-person shooter games [13]. Liaw et al. used GAs to evolve game agents that work as a team [14]. Othman et al. discuss using simulations to evolve an AI agent for tactical purposes [15]. Similar to the literature on optimization in aircraft engagement, these methods do not account for the cost of evaluating objective functions for optimization (e.g., number of human experiments or more generally financial/computational expense). They also do not attempt to construct a model of the true objective function (with uncertainty quantification) to enable better robustness and adaptation to the uncertainties involved with predicting and evaluating the outcome of real-time competitions between human/AI decision agents.

B. Formal Problem Definition

For simplicity and ease of developing the underlying theory, this work focuses on the problem of adjusting the parameters of a single AI decision maker competing against another AI decision maker in an aerial combat simulation. In this case, the “blue” agent represents the AI whose parameters are to be optimized, whereas the “red” agent represents an adversary whose decision-making parameters remain fixed (but are not necessarily known to blue). The red agent in this case can also be viewed as a proxy for a human pilot in a training scenario. The decision-making

Table 1 Agent behavior parameter elements $x_b(j)$

j	Abbreviation	Parameter name	Description
1	attackrange	AttackRange	Distance at which attacker begins active pursuit of detected adversary
2	endspeed	EndSpeed	Speed to use while on patrol or acquiring ground targets
3	fixedvertical	FixedVertical	Altitude to use while on patrol or acquiring ground targets
4	intspeed	InterceptSpeed	Speed to use while in attack posture approaching adversary
5	launch	LaunchDelay	Time to wait to fire weapon after lock has been achieved
6	maxaz	MaxAimuth	Maximum azimuth to allow adversary when in pursuit
7	maxel	MaxElevation	Maximum elevation to allow adversary when in pursuit
8	minagl	MinAgl	Minimum above-ground-level altitude when engaged in combat
9	select	SelectDelay	Delay in selecting weapon to employ
10	track	TrackDelay	Delay in responding to sensor indication of adversary track

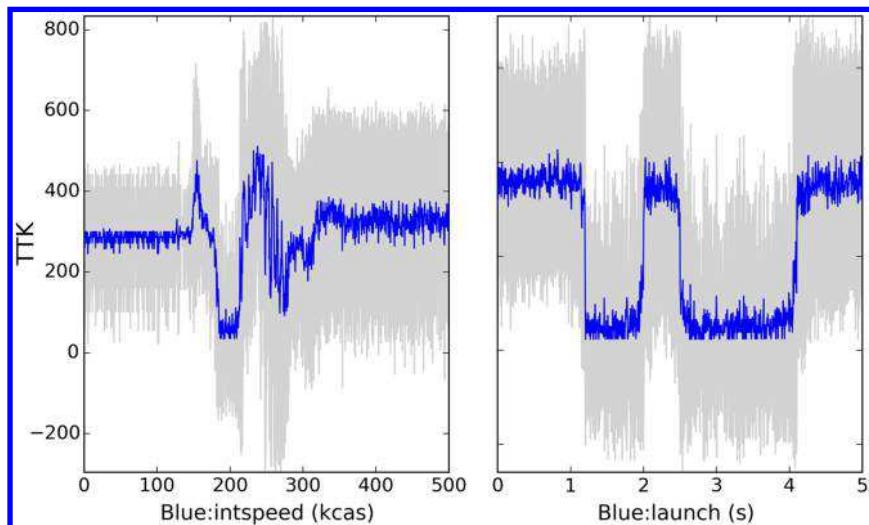
Table 2 Simulation objective metrics $y_m(\cdot)$

m	Abbreviation	Metric name	Description
1	survived	Survived	Boolean representing whether the agent survived or not
2	TST	TotalSimulatedTime	Duration of the simulation
3	numKills	NumberOfKills	Number of adversary kills achieved
4	MOC	MissionObjectivesCount	Number of ground targets acquired
5	cTTK	CumulativeTimeToKill	Cumulative time between weapons lock and target elimination
6	cTOO	CumulativeTimeOnOffense	Cumulative time in pursuit
7	cTOD	CumulativeTimeOnDefense	Cumulative being pursued
8	WU	WeaponsUsed	Number of weapons employed during the simulation
9	cTAL	CumulativeTimeToAchieveLock	Cumulative time between detection of adversary and weapons lock
10	EMM	EnergyManagementMetric [16]	Integrated total energy delta relative to adversary in the course of the simulation
11	SAM	StickAggressivenessMetric [16]	Measure of aggressiveness applied to control of the aircraft in the course of the simulation

parameter vectors are given by \mathbf{x}_r and \mathbf{x}_b for red and blue, respectively. For the aerial combat domain considered here, \mathbf{x}_r and \mathbf{x}_b are assumed to consist of variables such as those shown in Table 1. These parameters are taken from a proprietary aerial combat simulation AI developed by OrbitLogic, Inc. and define different threshold settings for triggering autonomous pilot behaviors defined by a state machine. These variables define the space over which optimization will occur in this paper, although other AI and decision-making parameters could also be used.

The goal is to minimize some desired engagement performance objective function $y_m(\mathbf{x}_r, \mathbf{x}_b)$ with respect to \mathbf{x}_b while \mathbf{x}_r is held constant. Here, $y_m(\cdot)$ must be evaluated using a high-fidelity combat simulation. The subscript m denotes that there may be multiple objective functions for a single objective simulation environment y . Table 2 lists some of the possible objective metrics for the simulation, which were compiled from the aforementioned literature on air combat training evaluations [16]. Each of the metrics describes a post hoc evaluation of blue's performance against red, based on telemetry data, control input data, and other relevant dynamic variables recorded during each engagement. Although each of the parameters and the outcome metrics are well defined, each simulation is expensive. From Table 1, we also see that the optimization can take place in a fairly high-dimensional space.

Figure 1 shows some examples of $y_m(\cdot)$ for typical simulation runs of the one-on-one engagement scenario defined later in Sec. III. These figures were produced by holding all \mathbf{x}_r parameters constant, as well as all \mathbf{x}_b parameters, except the one listed. Data were obtained by sampling the simulated engagement 1000 times along the horizontal axis and repeating that process 20 times. Clearly, these performance functions are highly noisy and volatile due to stochastic uncertainties. These uncertainties are due to environmental conditions such as, wind speed, terrain, and weather, as well as other sources such as error from sensors, psychomotor time delays, and initial placement of the fighters. Furthermore, it is clear that $y_m(\cdot)$ will generally not be a convex function of \mathbf{x}_b . Consequently, a specialized class of optimization is needed that can 1) handle inherently

**Fig. 1** Mean and 2σ bounds for TTK objective function vs x_b variables Blue:intspeed and Blue:launch.

noisy $y_m(\cdot)$ data, 2) identify the best local minima with respect to \mathbf{x}_b , and 3) reduce the number of evaluations needed to find the (expected) minimum of $y_m(\cdot)$. Bayesian optimization will now be introduced as a tool that can address these inherent challenges.

C. Bayesian Optimization

The goal of optimization is to minimize some objective function $y: \mathcal{X} \rightarrow \mathbb{R}$, where $\mathcal{X} \in \mathbb{R}^d$ is the search or solution space, and the element $\mathbf{x}^* \in \mathcal{X}$ is the minimizer, such that $y(\mathbf{x}^*) \leq y(\mathbf{x})$, $\forall \mathbf{x} \in \mathcal{X}$. Typically, the solution space is bounded for global optimization, where $\mathbf{x}(i) \in [x(i)_l, x(i)_u]$ for lower bound $x(i)_l$ and upper bound $x(i)_u$ for element i of \mathbf{x} . In classical optimization theory, one assumes the mathematical description of y as a function of \mathbf{x} is exactly known (e.g., $y = x^2 + 3x^3$ for scalar $\mathbf{x} = x$). In this case, it is fairly inexpensive/fast to evaluate y and apply standard optimization techniques. However, when the mapping from \mathbf{x} to y is not known explicitly, then optimization typically requires the evaluation of a “black-box” function. In many applications, such as the one considered here (i.e., where y is a high-fidelity simulator, the output of which is final engagement metrics), the black-box evaluations of y can be expensive and slow and can produce noisy results for the same input values \mathbf{x} .

The goal of Bayesian optimization is to find the minimizer of some unknown noisy objective function that is costly to evaluate, while also learning about it at the same time. It is so named because Bayesian inference, a key concept from probability theory, is applied during the optimization process. That is, the process uses an initial prior belief $p(y)$ over potential y functions, which can be updated by subsequent observations or evidence E consisting of sample evaluations of y for different values of \mathbf{x} . Mathematically, this leads to an application of Bayes’s rule: $p(y|E) \propto p(E|y)p(y)$ (see Appendix A.1 of [17] for more detail). The quantity $p(E|y)$ is also known as the observation likelihood, whereas $p(y|E)$ is the posterior (i.e., the updated probability of y after E is observed). Hence, given some initial prior belief about y and evidence E about the actual shape of y , an updated posterior belief about y can be formed. As long as both the prior and likelihood are consistent with the true nature of y , then the law of large numbers ensures that the posterior $p(y|E)$ converges with high probability to the true y in the limit of infinite observations E covering \mathcal{X} .

Bayesian optimization uses black-box point evaluations of y to find \mathbf{x}^* . This is accomplished by maintaining beliefs about how y behaves over all \mathbf{x} in the form of a “surrogate model” \mathcal{S} , which statistically approximates y and is easier to evaluate (again y might be an expensive high-fidelity simulation). During optimization, \mathcal{S} is used to determine where the next point sample evaluation of y should occur, to update beliefs over y and thus simultaneously improve \mathcal{S} while finding the (expected) minimum of y as quickly as possible. The key idea is that, as more observations are sampled at different \mathbf{x} locations, the \mathbf{x} samples themselves eventually converge to the expected minimizer \mathbf{x}^* of y . Because \mathcal{S} contains statistical information about the level of uncertainty in y [i.e., related to $p(y|E)$, the posterior belief], Bayesian optimization effectively leverages probabilistic explore-exploit behavior to learn an approximate model of y while also minimizing it. The main components of the Bayesian optimization process are 1) a surrogate model \mathcal{S} , which encodes statistical beliefs about y in light of previous observations and a prior belief, and 2) an acquisition function $a(\mathbf{x})$, which is used to intelligently guide the search for \mathbf{x}^* via \mathcal{S} .

1. Surrogate Model

The surrogate model \mathcal{S} must fulfill some basic properties to be useful in Bayesian optimization. First, it must approximate the function in areas where it has not yet been evaluated. Second, when it is evaluated, it must yield both a predicted value and a corresponding uncertainty to quantify the possibility that the optimum is located at some location \mathbf{x} . Several different kinds of surrogate models can be used, e.g., Gaussian processes (GPs) [17], random forests [18], and neural networks [19]. Different types of surrogate models are also available for categorical/integer inputs and for problems where \mathbf{x} is very high-dimensional [19,20]. GPs are by far the most common model used as a surrogate model in Bayesian optimization and are the focus of the work here. They are especially useful for continuous optimization problems of relatively low dimensionality ($d \approx 10$) [21]. Henceforth, the acronym GPBO refers to Bayesian optimization using a GP surrogate model \mathcal{S} .

A GP can be used to describe a distribution over functions; it is more formally defined as a collection of random variables, any finite number of which have a joint Gaussian distribution [see Eq. (1)] [17,22]. The process is completely specified by its mean function $m(\mathbf{x})$ [Eq. (2)] and its covariance function $k(\mathbf{x}, \mathbf{x}')$ [Eq. (3)]:

$$f(\mathbf{x}) \sim \mathcal{GP}(m(\mathbf{x}), k(\mathbf{x}, \mathbf{x}')) \quad (1)$$

$$m(\mathbf{x}) = \mathbb{E}[f(\mathbf{x})] \quad (2)$$

$$k(\mathbf{x}, \mathbf{x}') = \mathbb{E}[(f(\mathbf{x}) - m(\mathbf{x}))(f(\mathbf{x}') - m(\mathbf{x}'))] \quad (3)$$

In theory, $m(\mathbf{x})$ could be any function; as is common practice, this work assumes that m is zero for ease of exposition. The covariance function, also called a kernel, is a function that maps \mathbf{x} and \mathbf{x}' to a scalar, or $k: (\mathbf{x}, \mathbf{x}') \rightarrow \mathbb{R}$. These functions must be specified a priori and would ideally be based on some prior knowledge of the properties of y , such as its smoothness; the kernel used in our research is presented in Sec. III.E.

A kernel must be positive-semidefinite (PSD) to be valid. A kernel is said to be PSD if it produces a Gram matrix K , with individual elements $[K_{i,j}]$ given by $k(\mathbf{x}_i, \mathbf{x}_j)$ that is PSD given a set of training data $X = \{\mathbf{x}_1, \dots, \mathbf{x}_n\}$. Kernels are further discussed in Sec. III.E. $K(X, X)$ is the Gram matrix defined by kernel k :

$$K(X, X) = \begin{bmatrix} k(\mathbf{x}_1, \mathbf{x}_1) & k(\mathbf{x}_1, \mathbf{x}_2) & \dots & k(\mathbf{x}_1, \mathbf{x}_n) \\ k(\mathbf{x}_2, \mathbf{x}_1) & k(\mathbf{x}_2, \mathbf{x}_2) & \dots & k(\mathbf{x}_2, \mathbf{x}_n) \\ \vdots & \vdots & \ddots & \vdots \\ k(\mathbf{x}_n, \mathbf{x}_1) & k(\mathbf{x}_n, \mathbf{x}_2) & \dots & k(\mathbf{x}_n, \mathbf{x}_n) \end{bmatrix} \quad (4)$$

Given n training observations, the individual elements of the covariance matrix $K(X, X) \in \mathbb{R}^{n \times n}$ are the covariances $k(\mathbf{x}_i, \mathbf{x}_j)$ between \mathbf{x}_i and \mathbf{x}_j for all pairs of training data. For Gaussian observation likelihoods, the joint distribution of n training outputs $f(X) \in \mathbb{R}^{n \times 1}$ and p test outputs $f_*(X_*) \in \mathbb{R}^{p \times 1}$ for p test inputs $\{X_* = \mathbf{x}_{*1}, \dots, \mathbf{x}_{*p}\}$ can be written as

$$\begin{bmatrix} \mathbf{f} \\ \mathbf{f}_* \end{bmatrix} \sim \mathcal{N}\left(0, \begin{bmatrix} K(X, X) & K(X, X_*) \\ K(X_*, X) & K(X_*, X_*) \end{bmatrix}\right) \quad (5)$$

$$K(X_*, X) = \begin{bmatrix} k(x_{*1}, x_1) & k(x_{*1}, x_2) & \cdots & k(x_{*1}, x_n) \\ k(x_{*2}, x_1) & k(x_{*2}, x_2) & \cdots & k(x_{*2}, x_n) \\ \vdots & \vdots & \ddots & \vdots \\ k(x_{*p}, x_1) & k(x_{*p}, x_2) & \cdots & k(x_{*p}, x_n) \end{bmatrix} \quad (6)$$

Thus, given observed training values X and \mathbf{f} , predictions of \mathbf{f}_* can be made at new input locations X_* . For inference, the conditional values of the test outputs are of interest:

$$\mathbf{f}_* | X_*, X, \mathbf{f} \sim \mathcal{N}(\mu(X_*), \sigma^2(X_*)) \quad (7)$$

$$\mu(X_*) = K(X_*, X)K(X, X)^{-1}\mathbf{f} \quad (8)$$

$$\sigma^2(X_*) = K(X_*, X_*) - K(X_*, X)K(X, X)^{-1}K(X, X_*) \quad (9)$$

Here, $K(X, X_*) \in \mathbb{R}^{p \times n}$, so that $\mu(X_*) \in \mathbb{R}^{p \times 1}$ and $\sigma^2(X_*) \in \mathbb{R}^{p \times p}$. Equation (7) gives the expression of the conditional distribution of \mathbf{f}_* given test points X_* and training data X and \mathbf{f} . The mean and variance of this predictive distribution are found via Eqs. (8) and (9). In the context of Bayesian optimization, the surrogate model provides statistical information [i.e., mean and variance from Eqs. (8) and (9)] of how the underlying objective function y behaves for all possible values \mathbf{x}_* that have not yet been sampled.

2. Acquisition Function

Generally, an acquisition function operates on the surrogate model and yields a function whose optimum reflects the most promising location for the next function evaluation to occur. More formally, the acquisition function is defined as $a: (\mathbf{x}, \mathcal{S}) \rightarrow \mathbb{R}$. In Eq. (10), the general surrogate model \mathcal{S} has been replaced with a Gaussian process to reflect the model used for GPBO. To have more compact notation, we hereafter refer to the acquisition function as $a(\mathbf{x})$, which assumes the inclusion of the surrogate model as an argument:

$$a(\mathbf{x}) \triangleq a(\mathbf{x}, \mathcal{GP}(m(\mathbf{x}), (k(\mathbf{x}, \mathbf{x}')))) \quad (10)$$

Given an acquisition function, GPBO selects $\hat{\mathbf{x}} = \operatorname{argmax}_{\mathbf{x}} a(\mathbf{x})$ as the next location to be evaluated in the search process.

An important characteristic of any acquisition function is that it should not lead to greedy or myopic behavior. That is, $a(\mathbf{x})$ should enable sufficient exploration and modeling of y by sampling new locations \mathbf{x} that will improve \mathcal{S} . At the same time, $a(\mathbf{x})$ must efficiently exploit \mathcal{S} to reach the expected minimum of y as quickly as possible. There are many ways to define $a(\mathbf{x})$, although no single method is superior to all others in all applications (a manifestation of the “no free lunch” theorem) [23–25]. A few widely used methods are considered here: expected improvement (EI), upper confidence bound (UCB), and Thompson sampling (TS).

The EI method seeks to select the next sample point to maximize the statistically expected improvement in the optimum when the current best minimizer is \mathbf{x}^+ . The EI function (see [26]) is defined by

$$a_{\text{EI}}(\mathbf{x}) = \begin{cases} (\mu(\mathbf{x}) - f(\mathbf{x}^+))\Phi(Z) + \sigma(\mathbf{x})\phi(Z), & \sigma(\mathbf{x}) > 0 \\ 0, & \sigma(\mathbf{x}) = 0 \end{cases} \quad (11)$$

$$Z = \frac{\mu(\mathbf{x}) - f(\mathbf{x}^+)}{\sigma(\mathbf{x})} \quad (12)$$

where $\mu(\mathbf{x})$ is the mean predicted value of the GP at \mathbf{x} ; $\sigma(\mathbf{x})$ is the predicted standard deviation at \mathbf{x} ; $f(\mathbf{x}^+)$ is the best observed value of the objective function^{**}; and $\Phi(Z)$ and $\phi(Z)$ are the PDF and CDF of the standard normal distribution Z , respectively.

The UCB acquisition function was originally inspired by the UCB regret in multi-armed bandit problems [27]. The key idea is to bound regret in a sequential optimization process, where regret is defined as the difference between the actual strategy (based on decisions made with imperfect information) and the ideal strategy (decisions made with perfect information). UCB is defined as

$$a_{\text{UCB}}(\mathbf{x}) = \mu(\mathbf{x}) + \beta\sigma(\mathbf{x}) \quad (13)$$

where $\mu(\mathbf{x})$ and $\sigma(\mathbf{x})$ again come from the GP, and $\beta > 0$ is a hyperparameter. UCB thus simply takes the mean GP prediction and adds some multiple β of the standard deviation to the mean GP prediction for every \mathbf{x} . Note that, in minimization problems, the lower confidence bound obtained by subtracting $\beta\sigma(\mathbf{x})$ is used instead; to remain consistent with the predominant terminology in the GPBO literature, however, this acquisition function will still be referred to as the UCB for the remainder of this paper. One drawback for this function is that β must be tuned; although principled ways exist to do this (e.g., [23]), alternatives are often sought to avoid introducing this extra dependence.

**This should be substituted by the best mean prediction if the objective function is stochastic.

Algorithm 1 GPBO

```

1: Initialize surrogate model with seed data  $f$  and hyperparameters
2: while termination criteria not met do
3:    $x_j = \operatorname{argmax}_X a(\mathbf{x})$ 
4:   Evaluate  $y(x_j)$ 
5:   Add  $y(x_j)$  to  $f$ ,  $x_j$  to  $X$ , and update hyperparameters
6: end while

```

Thompson sampling (TS) involves drawing functions directly from the GP surrogate. In practice, sampling a function from Eq. (7) means calculating the predicted joint GP mean and covariance at a finite set of points in \mathbf{x}_* and then sequentially drawing f_* values at these points based on the resulting conditional GP means and variances. Once a function sample is obtained, its optimum is easily found, and the objective function is evaluated at that point. Because of its stochastic nature, TS naturally incorporates explore-exploit behavior. TS also naturally lends itself to parallelized searches and (from the law of large numbers) will tend to identify the most likely locations for \mathbf{x}^* for sufficiently large samples sizes. However, because the complexity of GP inference is cubic in n , sampling GPs with sufficient resolution for $d \geq 3$ is usually quite expensive [23].

3. Bayesian Optimization Procedure and Example

Algorithm 1 gives the GPBO procedure. The termination criteria could be based on iteration thresholds, tolerances on changes to the optimum \mathbf{x} and/or y between iterations, or other methods [28]. In practice, the performance of Bayesian optimization depends greatly on the selection and parameterization of the surrogate model, along with the number and placement of the initial training observations (i.e., seed points). It is also important to consider how samples are drawn with respect to the acquisition function. These issues are addressed in more detail for the aerial combat simulation application in Sec. III.

The statement on line 5 merits extra attention. Because GPBO is an iterative process, newly sampled data \mathbf{x}_i and $y(\mathbf{x}_i)$ are added to the GP surrogate model each iteration. The parameters of the covariance kernel are recalculated by methods discussed in Sec. III.E. As a result, the acquisition function changes from one iteration to the next (sometimes drastically); this behavior will be highlighted in the following example.

Figure 2 shows a few iterations of the GPBO process for an illustrative toy example based on the one-dimensional (1-D) Forrester function,

$$y_F = (6x - 2)^2 \sin(12x - 4) \quad \text{for } x = [0, 1] \quad (14)$$

which is a canonical optimization objective function [29] with multiple local maxima, multiple local minima, and a saddle point. In this example, Gaussian noise with zero mean and standard deviation equal to 0.5 has been added to y_F (indicated by the lightly shaded gray region).

The columns in Fig. 2 show successive iterations of the GPBO algorithm; the top and bottom rows show \mathcal{S} and $a(\mathbf{x})$, respectively. In this case, $n = 6$ seed points have been randomly selected for the GP surrogate \mathcal{S} . These points are shown in the top-left plot, along with the associated GP predictions. The values from the GP are obtained using Eq. (7) to evaluate a fixed mesh of points over the range of \mathbf{x} . Note that, with only six seed data points, the GP does not represent y_F very well.

The second row of Fig. 2 shows the acquisition function $a(\mathbf{x})$, which in this example is given by the UCB. On the first iteration, $a(\mathbf{x})$ is maximized at approximately $\mathbf{x} = 0.4$ (star); y_F is then evaluated at this location on the second iteration of GPBO, and a new data point is added to the GP (triangle), so that the sampling-update cycle repeats. Also, notice that the GP uncertainty increases and decreases in certain areas from iteration to iteration; this is due to the fact that the hyperparameters governing the covariance function [Eq. (3)] are updated when more data is acquired (in particular, the GP tends to become less uncertain in areas it has sampled already). In the four iterations shown here, although the GP does not represent y_F very accurately, it does manage to approximate y_F accurately in regions of interest (i.e., where the minimum might actually exist). On the fourth iteration, y_F is evaluated near the true minimize, and another point nearby this minimum will be selected via $a(\mathbf{x})$ for evaluation on the fifth iteration.

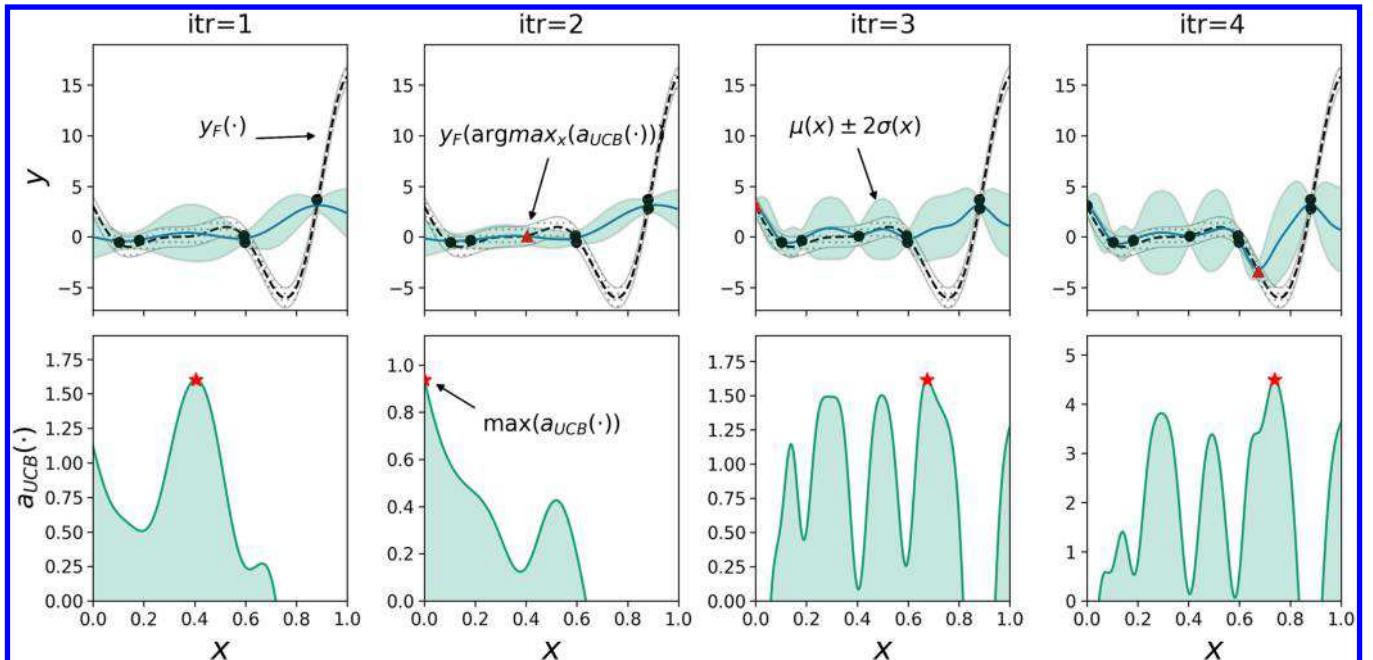


Fig. 2 GPBO on Forrester function (triangles: new data added to GP; stars: UCB acquisition function maximum).

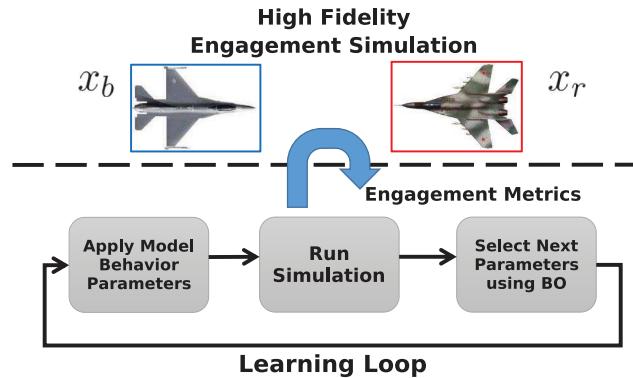


Fig. 3 Learning engine diagram.

This example highlights several key takeaways. First, the number and location of seed points affects the initial representation of the surrogate function and acquisition function. Second, the GP attempts to estimate the true objective function but only has information given by previous experiments. Third, the GP changes during optimization when trying to relearn the kernel hyperparameters with the newly acquired sample data. Finally, $a(\mathbf{x})$ will change when the GP surrogate \mathcal{S} changes. This means that, if the GP is a poor model of y and does not accurately represent uncertainty in unsampled parts of \mathcal{X} , then the GPBO process will tend to perform poorly. An ideal sampling strategy would yield a GP that provides a good global model of y with as few evaluation points as possible.

III. Bayesian Optimization for Decision-Making Artificial-Intelligence Adaptation

Bayesian optimization has emerged as a critical tool for “autotuning” various machine learning and AI algorithms [30–32] and has been applied to modeling of user preferences and reinforcement learning [24]. It is well suited for optimizing bounded objective functions that are unknown and expensive to evaluate and is often able to do this with the fewest function evaluations as compared to other competing methods [26]. These properties make it appealing for optimally adapting the behavioral parameters of an AI decision maker. In this section, the overall GPBO learning process for the aerial combat simulation application is described first. Then, specific strategies for coping with the major practical implementation issues for GPBO are described.

A. Learning Engine

For simplicity, the focus hereafter will be on adapting blue’s behavioral parameters only for simulated AI versus AI combat. As such, the red agent is treated as an AI agent with constant default parameter values \mathbf{x}_r (although the learning procedure does not require this assumption). Figure 3 depicts the process by which the optimum blue parameters \mathbf{x}_b are learned. The bottom portion of the figure shows the high-level learning loop, whereas the top of the figure represents the simulation environment. An initial set of agent parameters is chosen first, and then a simulation is run using the parameters for each agent using the Mission Simulation System (MSS) simulation environment developed by Orbit Logic, Inc. for live–virtual–constructive (LVC) applications. MSS employs very high-fidelity models of aircraft; their avionics system; installed sensors, weapons, and countermeasures; and external situation sensors. Additionally, AI agent models were developed to support tactics development, wargaming, and pilot training. The AI agents allow complex blue/red engagements to be played out without any humans in the loop, or with a mix of agents and humans. Table 1 in Sec. II.B shows the behavioral parameters for the proprietary pilot agent AI developed by Orbit Logic. The overall learning approach is agnostic to the internal details of the agent’s AI, and thus any other AI with tunable behavior parameters could be used instead.

At the end of a simulated engagement, one of the summary engagement metrics $y_m(\mathbf{x}_r, \mathbf{x}_b)$ shown in Table 2, such as time to kill (TTK), energy management, number of objectives completed, etc., is calculated based on recorded data for the red and blue aircraft and then delivered to the GPBO component. The GPBO algorithm selects one or more \mathbf{x}_b values to implement for the next round of engagement simulations that will potentially optimize the desired $y_m(\mathbf{x}_r, \mathbf{x}_b)$. This initiates the simulation–learning–parameter update cycle again; if more than one \mathbf{x}_b value is selected by GPBO, then each value results in a separate parallel instance of a follow-up simulated engagement. The cycle repeats until the desired termination condition is met. Possible termination criteria include 1) no significant changes in either \mathbf{x}_b or $y_m(\mathbf{x}_r, \mathbf{x}_b)$ are found, 2) the maximum number of optimization iterations has been reached, and 3) the maximum number of simulated engagements has been reached. Note that 2 and 3 are identical if only one \mathbf{x}_b value is selected by GPBO following a single engagement simulation instance.

B. Engagement Simulation Setup

A one-on-one engagement scenario was designed to collect relevant metrics for training and validating the GPBO tuning process. For ease of explanation and analysis, all results and discussion for GPBO in the air combat simulation application will be based on this single scenario. The scenario includes a single blue force fighter jet penetrating an adversary’s defended engagement zone. The simulation is run for $T_{\max} = 300$ s, or until either the blue or the red agent is eliminated. The blue agent’s primary objective is to engage in air-to-air combat against a red defending fighter jet to achieve theater control. The blue agent’s “goal” is to minimize the TTK, which is defined in Eq. (15). Thus, $\text{TTK} > 300$ represents blue being eliminated, $\text{TTK} < 300$ represents blue victory, and $\text{TTK} = 300$ denotes that both fighters survived:

$$\text{TTK} = \begin{cases} \text{SimTime,} & \text{if Blue survived} \\ 2 * T_{\max} - \text{SimTime,} & \text{otherwise} \end{cases} \quad (15)$$

Each agent is assigned a nominal flight plan that includes multiple waypoints that, when flown over, provide mission “scoring points”. The waypoints are arranged in such a way as to ensure that the fighters periodically encounter each other as illustrated in Fig. 4, instigating the employment of air-to-air combat logic. Engagement metrics are calculated at the conclusion of the simulation and then delivered to the learning engine.

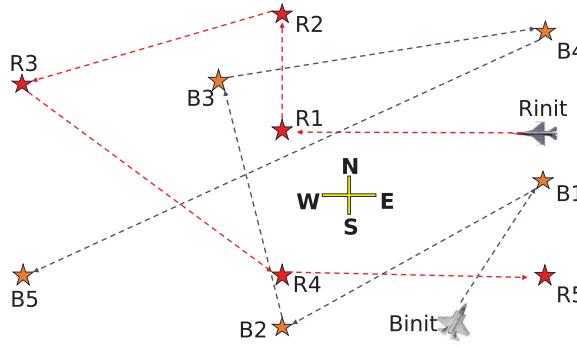


Fig. 4 Diagram depicting engagement space. Red and blue players begin at R_{init} and B_{init} .

The remaining subsections describe specific practical issues that arise for running the GPBO learning process and the solutions developed to address these.

C. Initial Data Seeding

For initialization, it is necessary to bootstrap the GP surrogate model with an adequate “seed” data sample. To get a good seed sample of the \mathbf{x}_b variable space, random sampling is often used [33]. However, to ensure sufficient coverage of the \mathbf{x}_b space, naïve uniform random sampling is inefficient because the samples are not correlated to one another and thus can be undesirably drawn in proximity to each other. Instead, a type of stratified quasi-random sampling called Latin hypercube sampling is used [34,35] to ensure a more uniform initial sampling of \mathbf{x}_b .

The success of GPBO rests heavily on the GP’s ability to globally model the true objective function. Without proper initial sample seeding, the GPBO algorithm can easily get trapped by local minima of the objective function and not escape. The number of random initial seed points to use mainly depends on the homogeneity of the objective function (i.e., whether it behaves similarly for different parts of the \mathbf{x}_b space) and the smoothness of the function. To the best of our knowledge, quasi-random sampling with a heuristically selected amount of seed data is currently state of the art for seeding [36,37]. It is common to use $n = 10d$ seed points, where d is the dimension of \mathcal{X} [26,38].

D. Selecting Sample Points from the Acquisition Function

The first step of the GPBO loop in Algorithm 1 is to find the arg max of the acquisition function $a(\mathbf{x})$. However, this is not a trivial matter. As shown in Fig. 2, it is very common for $a(\mathbf{x})$ to have many local maxima, regardless of which flavor of acquisition function is used. It is thus necessary to apply a global nonlinear/nonconvex optimization algorithm for this step. Although this means that another optimization algorithm must be run within GPBO, it is important to recall that this is, in fact, perfectly acceptable because the true objective function y_m is highly volatile and expensive to optimize using a global nonlinear/nonconvex optimization algorithm. As such, it is computationally far more efficient and operationally more acceptable/desirable to use such an optimization algorithm on $a(\mathbf{x})$.

This work makes use of the DIRECT (from “dividing rectangles”) optimization algorithm to find the best local maximum of $a(\mathbf{x})$. DIRECT is also widely used elsewhere in the GPBO literature [25,32,39,40]. This method uses the Lipschitz continuity properties of \mathcal{S} to bound function values in local rectangles and search accordingly [41]. Further details on DIRECT and a discussion of other possible global optimization algorithms for GPBO can be found in [23].

In standard GPBO, only a single new \mathbf{x}_b point is sampled according to $a(\mathbf{x})$ on each optimization iteration. However, as discussed in Sec. IV, it is also possible (and sometimes much more beneficial) to redefine $a(\mathbf{x})$ so that multiple \mathbf{x}_b points are sampled simultaneously on each iteration.

E. Updating the Kernel Hyperparameters

The kernel used to define k in this research is the Matérn kernel:

$$k_{\nu=3/2}(\mathbf{x}_{b,i}, \mathbf{x}_{b,j}) = \sigma_0 \left(1 + \frac{\sqrt{3}r_{ij}}{\ell} \right) \exp\left(-\frac{\sqrt{3}r_{ij}}{\ell}\right) \quad (16)$$

$$r_{ij} = \sqrt{(\mathbf{x}_{b,i} - \mathbf{x}_{b,j})^T (\mathbf{x}_{b,i} - \mathbf{x}_{b,j})} \quad (17)$$

with hyperparameters σ_0 and ℓ , which are the kernel amplitude and length scale, respectively. The Matérn kernel is widely used in GPBO because it has the useful property that different values of ν yield a kernel with different degrees of differentiability (where ν is nearly always taken to be half-integer to simplify the kernel expression). The Matérn kernel is guaranteed to be k times differentiable when $k \leq \nu$. Therefore, the kernel with $\nu = 3/2$ is $k = 1$ times differentiable and reflects the prior belief that the outcome engagement metrics y_i for the aerial combat simulations do not have higher-order derivatives. Note that Eq. (17) is the “isotropic” version of the Matérn kernel, where ℓ is scalar, and each dimension of r_{ij} has equal influence. Incorporating separate ℓ terms for each element of r_{ij} leads to the nonisotropic automatic relevance determination (ARD) version of the kernel. An ARD kernel can automatically learn to ignore parts of \mathbf{x}_b that are irrelevant to predicting the output [17] by scaling variables based on their contribution to the covariance function Eq. (3). However, ARD kernels lead to more hyperparameters than non-ARD kernels and thus can overfit data more easily. As is standard in GP regression, an additive observation noise variance σ_n^2 is also assumed for each training datum $f(\mathbf{x}_{b,i})$, where

$$f(\mathbf{x}_{b,i}) = y(\mathbf{x}_r, \mathbf{x}_{b,i}) + \epsilon_i \quad (18)$$

$$\epsilon_i \sim \mathcal{N}(0, \sigma_n^2) \quad (19)$$

Hence, the full set of hyperparameters $\Theta = \{\sigma_n^2, \sigma_0, \ell\}$ governs the behavior of the covariance function of the GP in Eq. (3).

Because the best Θ values are unknown a priori, they must be learned and updated during GPBO. Point estimation strategies based on maximum-likelihood estimation (MLE) are the most widely used in the GPBO literature for supervised learning of Θ . MLE maximizes the GP likelihood function $p(f|X, \Theta, k)$ with respect to Θ , where

$$\begin{aligned} p(f|X, \Theta, k) &= \mathcal{N}(\mu(X), K_f(\Theta, k)) \\ &= (2\pi)^{-D/2} |K_f(\Theta, k)|^{-1/2} \exp\left(-\frac{1}{2}(f - \mu(X))^T K_f(\Theta, k)^{-1}(f - \mu(X))\right) \end{aligned} \quad (20)$$

$$K_f(\Theta, k) = K(X) + \sigma_n^2 I \quad (21)$$

Fast gradient-based convex optimization techniques are most commonly used to minimize the negative logarithm of Eq. (20) (i.e., the negative log likelihood) because the required derivatives can be obtained analytically. However, because the GP likelihood is generally nonconvex, numerical optimization can converge to many different local optima for Θ . To further complicate matters, the best local optimum may be undesirable for learning with sparse data early on in the GPBO process because the associated Θ values typically lead to overfitting of the training data [17, 42]. This behavior is especially important to consider when trying to minimize the number of simulations for GPBO. Cross-validation (CV) methods are commonly used with MLE to avoid overfitting. In CV, the GP is trained on a designated training data set using MLE and then validated by assessing the predictive error on a completely independent test data set to ensure that the learned model is not overfit. Although unbiased, CV tends to produce undesirably high error variances for small data sets [43], making it difficult to reliably learn Θ .

Alternatively, maximum a posteriori (MAP) estimation can be used to avoid overfitting by adding hyperpriors $p(\Theta|k)$ on Θ . In this case, the posterior distribution for Θ from Bayes's rule,

$$p(\Theta|f, X, k) \propto p(f|X, \Theta, k)p(\Theta|k) \quad (22)$$

is maximized via gradient ascent. Considering the logarithm of Eq. (22), the $p(\Theta|k)$ hyperpriors act as additional “regularization” variables for the GP data log-likelihood function in Eq. (20). This can help smooth away poor local optima and prevent point estimates of Θ from overfitting or changing too drastically given small amounts of data. As the number of sample points added to the GP increases, the contribution of $p(\Theta|k)$ to the point estimate of Θ diminishes. In this work, $p(\Theta|k)$ hyperpriors are simply defined as uniform distributions for a broad range of Θ values. Following MAP updates to the hyperparameters, any predictions made through the surrogate function can be modified to account for hyperparameter uncertainty by integrating with respect to Θ (i.e., by marginalizing out the hyperparameters):

$$p(f_*|X_*, f, X, k) \propto \int_{-\infty}^{\infty} p(f_*|X_*, f, X, \Theta, k)p(\Theta|f, X, k) d\Theta \quad (23)$$

Because both $p(\Theta|f, X, k)$ and this integral are in general analytically intractable, approximation schemes like the Laplace method can be used. In the Laplace method, $p(\Theta|f, X, k)$ is approximated as a multivariate normal $\mathcal{N}(\hat{\Theta}_{\text{MAP}}, \hat{\Sigma})$, where $\hat{\Theta}_{\text{MAP}}$ is the MAP estimate, and $\hat{\Sigma}$ is the negative inverse Hessian of the RHS of Eq. (22) evaluated at $\hat{\Theta}_{\text{MAP}}$. Equation (23) then becomes analytically tractable, so that it can be used in the acquisition function. This approach effectively inflates the prediction uncertainty of the GP by the uncertainty in Θ , which is especially useful for the initial phases of GPBO when sparse data are available.

The specification of hyperpriors is an excellent place to build in expert knowledge of the system if it is available. For example, one might tightly constrain the hyperparameter governing the length scale, if there were some reason to believe that the length scale should be near a certain value. In the air combat scenario, where we did not have much a priori knowledge, we used uniform, weak, prior distributions to allow the observations to have more effect on the learning process.

IV. Serial, Parallel, and Hybrid Sampling Strategies

One of the strengths of GPBO is that it can greatly reduce the number of required evaluations of an expensive objective function, which in this case requires a full simulation of an aerial engagement. In fact, GPBO is often able to find the optimum in fewer iterations than any other known optimization method [26]. However, if multiple new simulations can be specified and possibly run in parallel, the optimization process can be greatly accelerated and improved in terms of reliability and repeatability (a very desirable characteristic for black-box optimization with noisy objective functions). In the context of air combat simulation training, such parallelization is especially attractive for adapting AI-versus-AI engagements as well as developing aggregated pilot performance evaluations (i.e., where instances of the same AI are simultaneously pitted one-on-one against different human-controlled red agents in different simulation instances).

Three approaches can be considered for parallel sampling-based search. The first approach is to use $a(\mathbf{x})$ to identify several different \mathbf{x}_b locations simultaneously (e.g., the best local optimum and other “nearly optimal” locations simultaneously) or a sequence of hypothetically optimal locations. This will be referred to as multipoint sampling (MS).^{††} The second is to sample the same maximum of $a(\mathbf{x})$ multiple times. This method will be referred to as repeat sampling (RS). The traditional, nonparallel approach to sampling new \mathbf{x}_b via maximization of $a(\mathbf{x})$ is denoted single sampling (SS), which is a special case where RS = MS = 1.

Figure 5 illustrates the SS, RS, and MS sampling methods on a much noisier version of the 1-D Forrester function shown earlier, using the parallel sampling acquisition functions described later for each approach. From this simple example, it is easy to see that (in just four optimization iterations) the resulting GP surrogate models for both the RS and MS methods do a much better job at estimating the statistics of the underlying objective function than the traditional SS approach. Clearly, this improvement is due in large part to the fact that both MS and RS benefit from additional function evaluations and thus can use more data in their GPs to make better predictions. However, note that the MS approach at iteration 2 is already significantly better than SS at iteration 4 (i.e., even after the same number of new data samples have been added to both GPs). Hence, strategic placement of additional samples earlier on in the GPBO process can yield significant benefits. As shown later in Sec. V, the improved performance for RS and MS over SS (given the same total number of function evaluations) also generally appears to hold in more complex problems. With this in mind, different techniques for implementing MS and RS methods are considered next.

^{††}Also known as batch sampling.

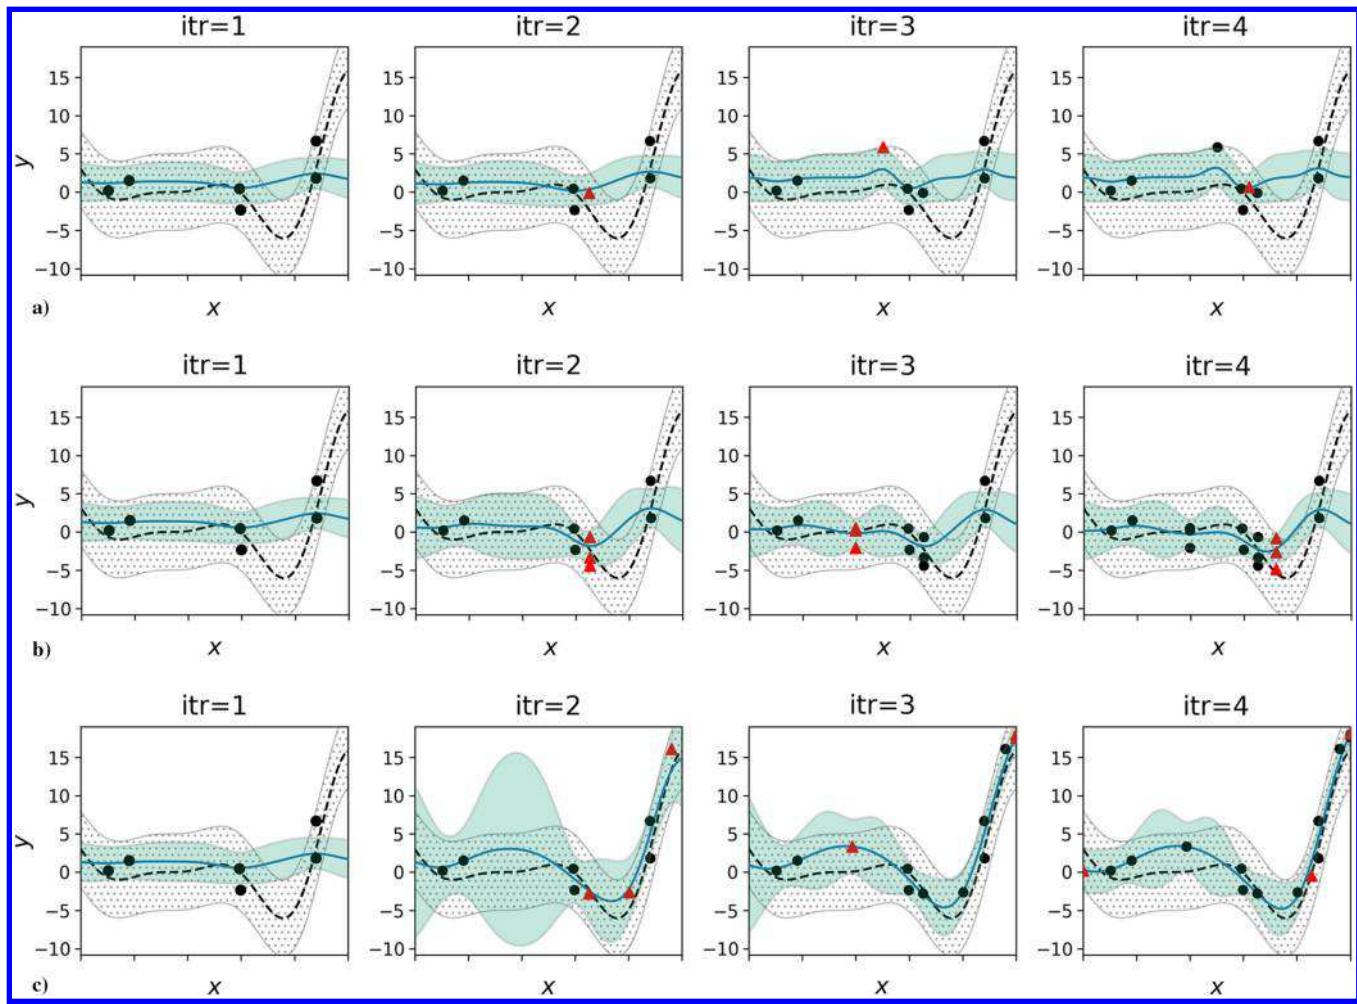


Fig. 5 GPBO comparison with UCB/variants: a) SS, b) RS = 3 and MS = 1, and c) RS = 1 and MS = 3 (GP-UCB-PE).

A. Multipoint Sampling

Multipoint sampling (MS) uses any one of several strategies to forecast which points might be of the most interest in future evaluations [31,37,44]. Three different methods are considered here.

The first q -EI method, introduced by [35,45], is similar to EI for traditional SS, except that, instead of optimizing over \mathbf{x}_b , the EI function is optimized over multiple ($q > 1$) sample locations \mathbf{x}_b at once. This method is notoriously expensive to evaluate because the dimensionality of $a(\mathbf{x})$ increases (e.g., with $q = 5$ and $d = 2$ leads to a 10-dimensional acquisition function), thus causing both calculation and optimization to become much more expensive. As an alternative, the approximation method introduced by [37] is often used for efficient computation of the q points. Their result bypasses the need for Monte Carlo simulations by applying the so-called “Tallis formula”, which uses a tractable method for calculating the conditional expectations of the multinormal distribution.

The second GP-UCB-PE method [44] uses the UCB function and a pure exploration (PE) technique to perform parallel search for the optimum. This approach is based on a simple premise: greedily select q points in sequence using the UCB acquisition function. This is done by first selecting the maximum of the UCB function as \mathbf{x}_1 , exactly as in a conventional SS approach. Next, taking advantage of the fact that the surrogate model conditional covariance predictions in Eq. (3) are independent of the exact value of f , the UCB can be recalculated, assuming that the proposed \mathbf{x}_1 is given as part of the covariance function, which allows for selection of \mathbf{x}_2 at the next UCB maximum. This process is repeated until q points have been selected.

Finally, a simple extension of basic Thompson sampling can also be used for MS. Instead of drawing a single function from the GP surrogate model and selecting its optimum, q functions are drawn together at once. Then, the optimum of each sample function is found, and experiments are run at each of the corresponding locations of an optimum. As in the case of Thompson sampling for classical SS, this approach is asymptotically efficient but still suffers from the curse of dimensionality.

B. Repeat Sampling

Repeat sampling (RS) has not been previously studied in the GPBO literature and is thus a novel contribution of this work. It is argued here that RS can be especially useful in problems with very noisy and volatile objective functions (i.e., those found in simulated aerial combat). The rationale behind using RS is intuitively related to the concept of “repeated measures” in experimental design [46]. Loosely speaking, if the variance for the outcome of any given individual experiment (simulation) is large, then repeating the experiment at the same setting for the independent variables (\mathbf{x}_b , in this case) allows statistical inference to be made with fewer experiments by controlling the variance and reducing sensitivity to outliers.

The connection to the mechanics of GPBO stems from the kernel function hyperparameter updates at the end of each optimization iteration. As discussed previously, the hyperparameter update requires solving a nonconvex optimization problem involving the GP data log-likelihood. If only a single noisy sample is evaluated in an iteration, the kernel parameter update could respond poorly to noise or outliers and thus converge to a poor local optimum in Θ space and/or produce an incorrect surrogate GP that overfits data. The resulting GP model S will then not accurately reflect the

true statistics of y_m , which in turn will corrupt $a(\mathbf{x})$ and can subsequently lead to suboptimal selections of $\hat{\mathbf{x}}_b$ for all remaining GPBO iterations. This issue is especially critical for the first few iterations of GPBO because the GP surrogate will be based on relatively small data sets. However, if a larger “statistically representative” sample of y_m is obtained for any given \mathbf{x}_b using RS in such cases, it becomes possible to reduce the sensitivity of the hyperparameter updates to noisy sample data. Specifically, RS makes it possible to improve estimation of the σ_n^2 noise hyperparameter in Eq. (20). This in turn can help improve estimation of the other hyperparameters in Θ early on in GPBO. Recall that because σ_n^2 characterizes the external noise term e_i in Eqs. (18) and (19), this hyperparameter effectively accounts for errors that are independent of \mathbf{x}_b in the GP surrogate model (i.e., errors in modeling y_m that can arise due to an imperfect choice of kernel function or kernel hyperparameters).

In theory, there are at least two ways to incorporate RS experiments into a GP. The first is by incorporating the summary value of multiple experiments (i.e., the mean observed value from n experiments). The second is to include all experiments individually. On the surface, the first method seems attractive because the computational complexity of a GP is cubic with respect to the number of data points. In practice, this approach does not yield desired results; instead, the second method must be used. Put simply, this is due to the fact that a summary observation does not directly encode the observed variance. In contrast, using all data fundamentally alters \mathcal{S} and makes the hyperparameter learning process more effective. This point will be discussed further in the next section.

In this work, acquisition functions for RS mirror those for the conventional SS approach, with the main difference being that multiple samples (instead of just one) are taken at the maximum of the acquisition function. Formal adaptations of MS acquisition functions could also be considered to constrain all samples to come from the same \mathbf{x}_b on a given optimization iteration.

C. Comparison of Multipoint Sampling and Repeat Sampling Effectiveness

It is posited here that MS and RS can significantly improve the surrogate model’s ability to capture underlying statistics (mean and covariance functions) of an underlying stochastic objective function, particularly early on in the GPBO process. To gain some insight into the expected gains offered by these approaches, it is first necessary to define a notion of statistical consistency or “goodness of fit” for the surrogate model GP. Suppose that a “ground truth” GP $\bar{\mathcal{S}}$ exists to describe the stochastic objective function, which has some known kernel function k and true hyperparameter values $\bar{\Theta}$. When performing GPBO, another surrogate model \mathcal{S} with the same kernel function k learns an estimate of the hyperparameters Θ from noisy data sampled from $\bar{\mathcal{S}}$. Because \mathcal{S} and $\bar{\mathcal{S}}$ use the same k , \mathcal{S} is statistically consistent with $\bar{\mathcal{S}}$ if $\Theta \rightarrow \bar{\Theta}$ as $n \rightarrow \infty$. Thus, the impact of using MS or RS can be understood (in this restricted regression scenario) by examining Eq. (20) following corresponding sample updates.

Figure 6 shows results from a toy problem that examines this consistency idea for MS and RS search strategies, when only small data sets are available for initial seeding. A 1-D “truth model” GP $\bar{\mathcal{S}}$ with known kernel hyperparameters $\bar{\Theta}$ (red vertical lines) was used to generate noisy objective function evaluation data. Fifteen random data points were then used to construct an approximate “initial seed” GP surrogate model \mathcal{S} using the same kernel function, but with an estimated Θ . The same number of new sample points were then added to \mathcal{S} using either MS (with GP-UCB-PE) or RS (using UCB). The resulting GP data likelihoods from Eq. (20) were then evaluated following a single update over a mesh of hyperparameter values (varying one element of Θ at a time, while holding the other hyperparameters fixed to their initial values); the rows show results for RS = MS = 1, RS = MS = 3, and RS = MS = 5, respectively, whereas the columns show variations of Eq. (20) with respect to each element of Θ for the Matérn kernel. The regression procedure was repeated 100 times with random initial seedings for each search approach, and likelihood curves are plotted for instances achieving median maximum-likelihood values.

The top row shows that both methods perform identically when RS = MS = 1, as expected. The remaining results show that RS and MS both perform comparably on the ℓ and σ_0 parameters and tend to remain biased early on in the GPBO process when the other hyperparameters are fixed. For σ_n^2 , MS has larger variance but a clear maximum far from the true σ_n^2 value (i.e., it is biased toward a larger σ_n^2 but has an appropriate amount of uncertainty to account for this). On the other hand, RS has a sharper local peak at the true σ_n^2 and more local optima, including a larger peak at a much larger σ_n^2 value. However, the local peak near the true value gets larger as the number of samples increases for RS, whereas MS reduces its bias more slowly with additional samples. This demonstrates the idea that, during hyperparameter optimization early on in the GPBO process, RS has a better chance of identifying the σ_n^2 value independently of the other hyperparameters, as new data points are sampled (i.e., as long as the best local optimum is found via MLE or MAP). It is also worth noting in this example that MS produces a more reliable and conservative estimate of the error uncertainty than RS, which tends to be more optimistic (as indicated the curvature of the likelihood functions). It is important to ensure that the hyperparameter error uncertainty is also correct because the GP data likelihood is recursively updated by GPBO in future sampling steps and thus must correctly steer the behavior of the surrogate and acquisition functions via hyperparameter updates.

In practice, as alluded to in the previous section, this means that including a summary statistic is indistinguishable from a single experiment; although the mean observed value will be less responsive to outliers, it will not convey information about the variance of the data. In contrast, using RS encodes information about the variance of the experiments into the kernel and fundamentally changes the likelihood of the hyperparameters.

A quick note on computational complexity is merited because it is a serious concern for GPs when data sets are large. There are methods that can make a GP “sparse” by distilling the data to fewer more important parts; this is a good idea when the GP is static but not when it is still being learned. In other words, if the GP is done being trained (i.e., the hyperparameters have been sufficiently learned), then the data can be made sparse to improve the speed of calculations. However, if the GP is still being learned, attempting to make the kernel sparse will modify the hyperparameter likelihood and will hinder the learning process. It is therefore recommended that, when learning hyperparameters (especially when very little data are available, as in an expensive optimization problem like this one), the GP be left “dense” and only reduced after the hyperparameters are no longer being learned.

D. Hybrid Repeat/Multipoint Sampling

The preceding considerations suggest that some combination of RS and MS might yield the “best of both worlds” for GPBO, which potentially enables more efficient and correct statistical modeling of noisy objective functions. Specifically, RS should be used to improve handling of sparse noisy samples early on in the process (i.e., to “anchor” the GP surrogate), whereas MS should be used to explore and exploit the GP surrogate more efficiently. This leads to a novel third sampling strategy, referred to here as hybrid repeat/multipoint sampling (HRMS). In this approach, the MS number allows GPBO to vary the total number of locations to investigate at each iteration, whereas the RS number allows GPBO to vary how many times to sample each location at each iteration. The remainder of the paper uses the convention that MS = m means m locations in \mathbf{x}_b space will be simulated in parallel, and RS = l means that the same location in \mathbf{x}_b space will be sampled l times.

In general, both m and l could vary over the course of the GPBO process, although compatible acquisition functions should be used for both the MS and RS portions. In this work, it is assumed that m and l are selected and fixed a priori for HRMS GPBO. For $m > 1$ and any $l \geq 1$, an MS acquisition function is used to select m distinct sampling locations, at which l repeated samples are evaluated. Because of the highly application-dependent nature of the problem, it remains a challenging and open research problem to determine how exactly m and l should be varied to achieve optimal GPBO performance. For instance, these parameters could be selected to minimize time to find \mathbf{x}_b^* in the fewest iterations possible. Alternatively, they could be selected to minimize the total number of objective function evaluations. The experimental results presented in Sec. V will provide some basic insights into this question in the context of the aerial combat simulation application.

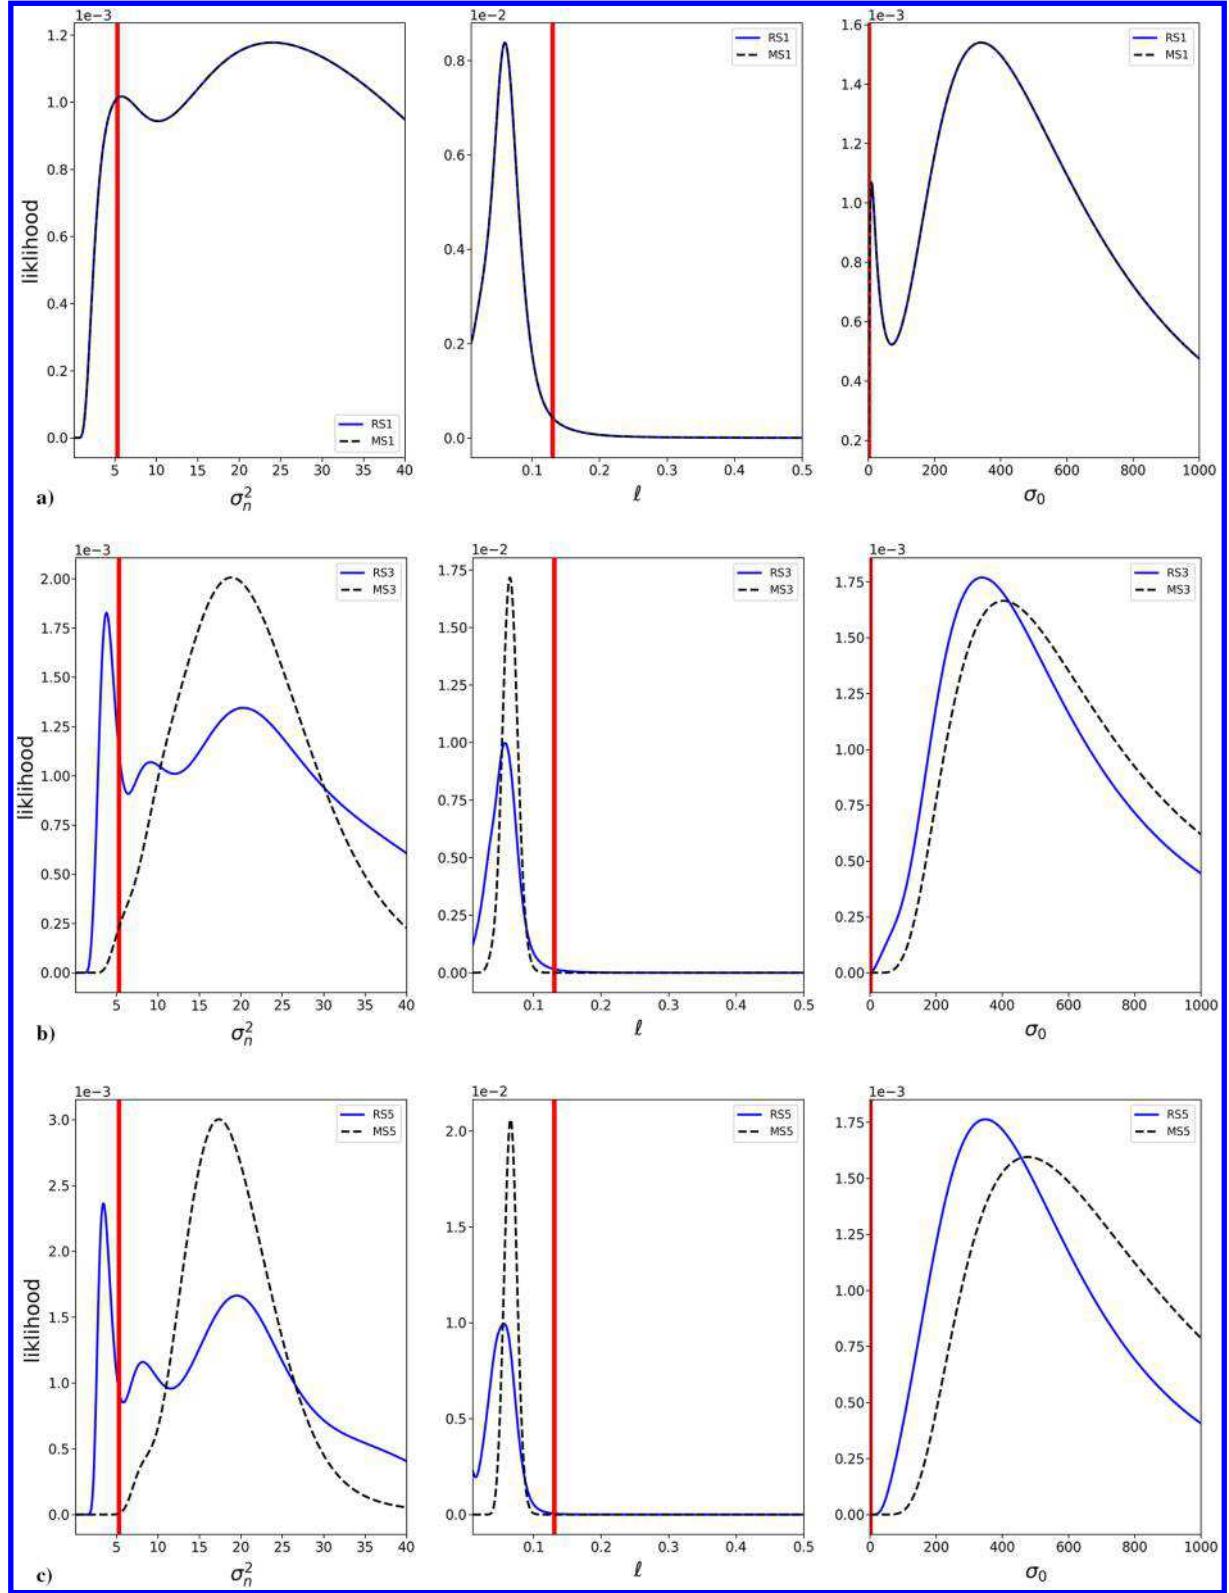


Fig. 6 RS vs MS effects on sparse data hyperparameter learning: a) RS = MS = 1, b) RS = MS = 3, and c) RS = MS = 5. Vertical lines: true values.

V. Simulation Results

This section presents the results of simulation studies that were conducted to investigate the performance of the different GPBO acquisition functions and sampling strategies on two different test problems. First, evaluation results for a standard benchmark optimization problem are presented, primarily to validate the newly proposed RS and HRMS sampling techniques on a nontrivial test case with a known solution. Then, results for the simulated air combat decision-making AI optimization problem described in Sec. III are presented, along with an analysis and discussion of the outcomes of GPBO on the AI pilot parameter adaptation process.

Three different acquisition functions were evaluated in each test case using SS, MS, RS, and HRMS sampling techniques. For SS and RS, these were expected improvement (EI) [26], upper confidence bound (GP-UCB) [27], and Thompson sampling (TS) [47]; for MS and HRMS, these

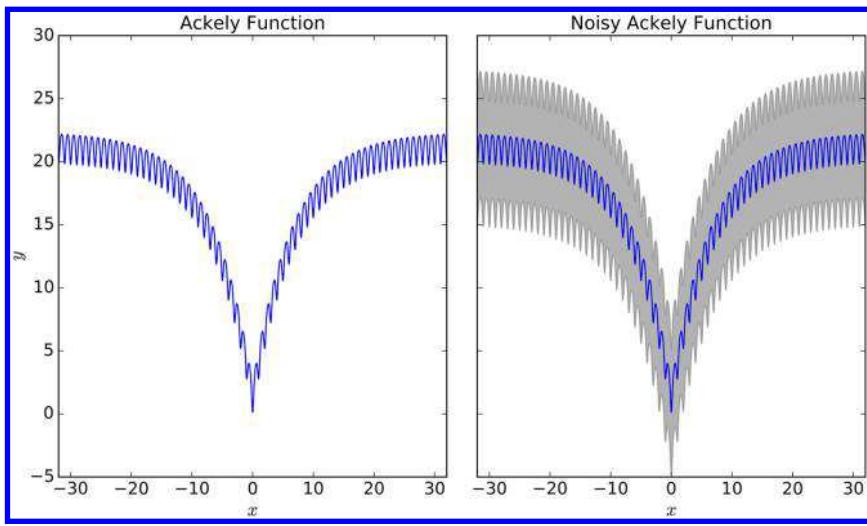


Fig. 7 Representations of standard Ackley function (left), and noisy Ackley function (right, with added variance $\sigma_y^2 = 25$).

were the corresponding batch sampling forms: q-EI [48], GP-UCB-PE [44], and multidraw TS. The different levels of RS and MS used are $RS = \{1, 3, 5, 10\}$ and $MS = \{1, 3, 5\}$ (where SS corresponds to $RS = MS = 1$, and HRMS corresponds to $RS > 1$ and $MS > 1$). The GPML toolbox [49] is used for GP representation and hyperparameter inference.

A. Ackley Three-Dimensional Function Experiments

To validate the usefulness of RS and HRMS, it is useful to apply these sampling methods to a standard benchmark optimization problem with a known minimum. The benchmark used here is a noisy version of the Ackley function [50], which is defined in n dimensions and has its largest local minimum value of 0 at $\mathbf{x} = 0$, where \mathbf{x} is an n -dimensional input vector. Figure 7 shows the standard and noisy versions of the Ackley function for $n = 1$. The deterministic base Ackley function has multiple periodic local minima that are closely spaced, and so this function is highly nontrivial to optimize using standard numerical optimization techniques. The noisy version of the Ackley function thus serves as an even more interesting and challenging benchmark for assessing GPBO techniques.

Although GPBO techniques are theoretically able to cope with the presence of multiple local expected minima, a poorly tuned or selected GPBO algorithm (like any other optimization technique) can easily get stuck at a local minimum and fail to explore the solution space. These issues become particularly challenging in high-dimensional settings, where sensitivity increases to initial seed conditions and poor local minima in the GP log-likelihood or log-posterior. Hence, this benchmark problem can provide some insight into how well different GPBO strategies can explore complex stochastic objective functions as well as accurately and repeatably represent them through the surrogate function.

The noisy Ackley function was optimized on 10 independent runs for each RS/MS GPBO configuration with $n = 3$ and $\sigma_y^2 = 25$, using 20 random initial seed points in the ranges $x_1 = (-32.768, 32.768)$, $x_2 = (-12.21, 32.768)$, and $x_3 = (-32.768, 5.14)$ for each dimension of the search space (thus bounding the largest local minimum for this problem at $\mathbf{x} = 0$). MAP hyperparameter estimation was used with the Laplace approximation, with Matérn 3/2 kernel hyperpriors set to uniform distributions between e^{-1} to e^2 .

Figure 8 is a convergence plot that illustrates the effect of HRMS on this problem where three different HRMS sampling strategies are applied using a time limit of 1 h (the number of function evaluations and iterations are not equal due to the complexity of different strategies). The main result of interest is found by comparing the right-hand side of each subplot. For example, the 03/03 approach yielded a better minimum than that of the 01/01 approach; it also had less variance between independent optimization attempts. This illustrates that there is utility in using HRMS on a function with high noise. This finding becomes even more critical when considering that, in true black-box optimization, there is no way to know when the algorithm converges, and thus using an HRMS approach (instead of SS) yields a more reliable and more accurate solution.

It should be noted, however, that to make a fair comparison between the different GPBO methods, the inherently different baseline computational costs of each approach need to be taken into account. Although optimization methods are typically compared by number of iterations until convergence, this comparison does not translate to the current comparative setup for GPBO because some approaches might use many more function evaluations per iteration (i.e., $RS = 10$, $MS = 5$ will use 50 function evaluations per optimization iteration,

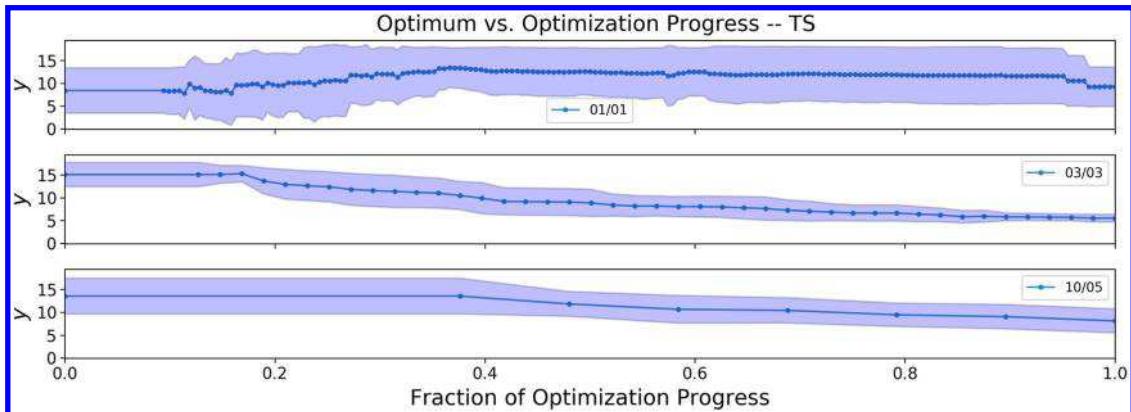


Fig. 8 Convergence for three RS/MS configurations (top: $RS = 1, MS = 1$; mean and 1σ shown for 10 experiments).

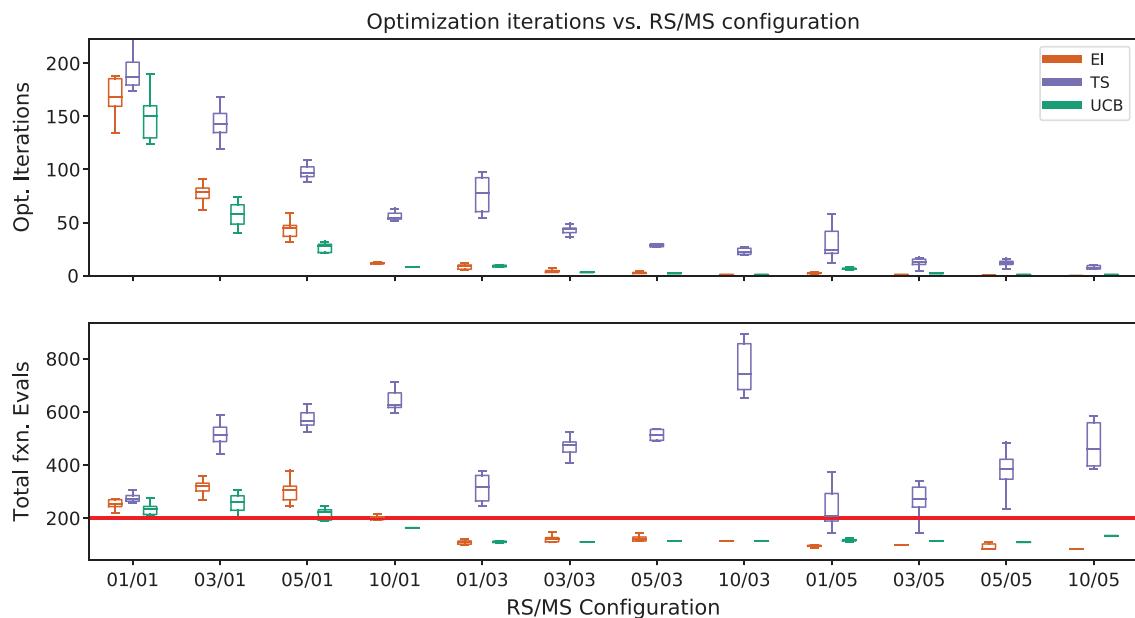


Fig. 9 Optimization iterations and total function evaluations for EI, TS, UCB acquisition functions, and HRMS configuration.

whereas $RS = 1$, $MS = 1$ will only use 1). Instead, the different GPBO approaches are more appropriately compared based on a fixed total number of function evaluations; also, because the real interest is the final result of the optimization, further figures will reflect the final results and not show typical convergence plots.

Figure 9 illustrates the difference between optimization iterations and total function evaluations across RS/MS configurations and acquisition functions. These results were generated after running optimization for 1 h, to highlight the different speeds at which each GPBO method runs. Generally, SS was the fastest approach yielding more function evaluations. Because of the expense of calculating the MS locations, the total number iterations for MS decreased as more evaluations were added per iteration. Notice that UCB had much fewer evaluations as MS increased, due to the high expense of the GP-UCB-PE algorithm. On the other hand, the cost of drawing MS random functions when using TS was relatively inexpensive (given that the number of sampled points is “small enough”). Note that the solid line on the bottom plot shows the “cutoff” point at or below which the estimates were taken to make Fig. 10, so that the different methods could be fairly compared for a given number of function evaluations.

The results of the optimization trials are summarized in Figs. 10 and 11, in terms of the outcomes of the estimates for each of the inputs and the optimum value for a fixed set of 200 function evaluations per method (or fewer if more could not be finished within the 1 h limit). Note that only the results for the Thompson sampling acquisition function are shown here; the results for the other acquisition functions show similar trends but are not displayed here.

The areas highlighted by blue boxes show the standard approaches using SS and MS. The orange box represents one example of the outcome of optimization using HRMS. Notice how the RS and HRMS configurations, in general, exhibit smaller variance and closer proximity to the true value. From this figure, we can conclude that there are RS/MS configurations that yield more reliable and accurate optimization results than those of SS or pure MS.

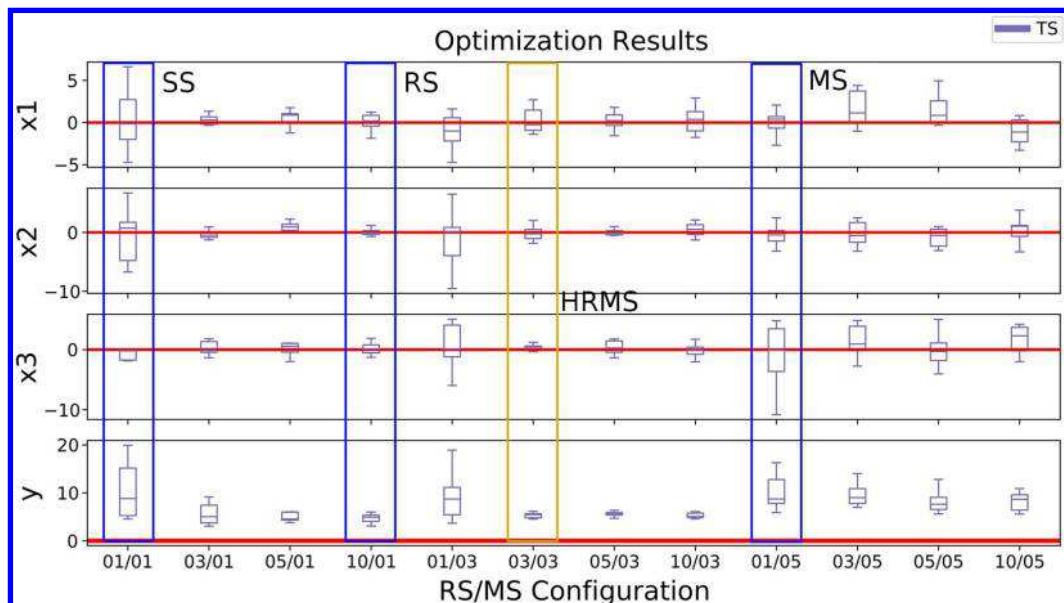


Fig. 10 TS-estimated x_1 , x_2 , x_3 , and y values by HRMS configuration (horizontal line: truth value).

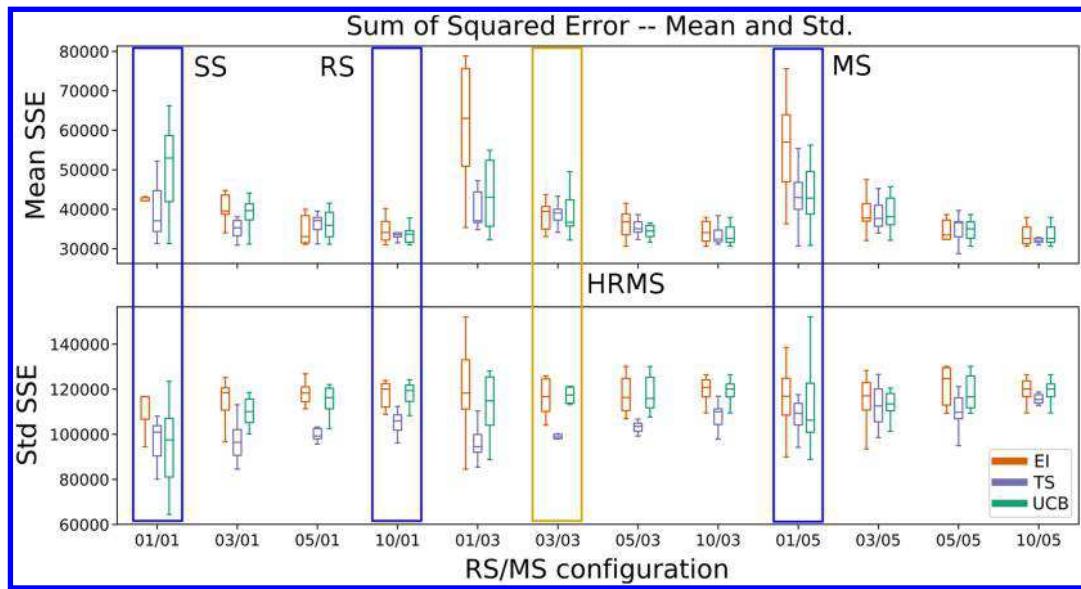


Fig. 11 SSE of GP mean (top) and σ (bottom) vs configuration (orange: HRMS example).

In Fig. 11, the sum of the squared errors (SSE, also called the residual sum of squares) between the final surrogate GP model and a ground truth objective model is shown. The SSE is calculated using

$$\text{SSE} = \sum_{i=1}^n (v_{i,\text{predicted}} - v_{i,\text{true}})^2 \quad (24)$$

where v is the value of interest at location i (i.e., the mean or variance). In other words, the SSE is the sum over n locations of the squared error between the predicted and true values. Here, the ground truth model is another higher-fidelity GP fit with 10,000 training samples drawn from the noisy three-dimensional (3-D) Ackley function; this allows for consistent comparison of the GP surrogate model mean and variance functions for the full range of input parameters.

A smaller SSE denotes a closer fit to the true objective function; an SSE with smaller variance reflects better repeatability. Notice the universal improvement of the RS and HRMS methods over the standard SS and MS configurations in modeling the objective function mean. The SSE of the standard deviation reduces variance with the HRMS method but does not seem to get closer to that of the true function. Although this behavior is not yet fully understood, one hypothesis is that this may be due to the sparsity of the experimentation using GPBO as well as the fact that there is higher uncertainty in areas that have not been explored. In other words, the mean representation is better, but the variance is still large in unexplored areas, similar to what is seen in iteration 4 of Fig. 2.

Some of the main insights from these results are summarized as follows.

- 1) Although it is possible to obtain reasonable optimization results using the standard SS and MS strategies, the solutions are highly sensitive and have high variance (as seen in Figs. 8 and 10).
- 2) Using HRMS, it is possible to not only outperform the accuracy of the predictions from SS and MS sampling approaches but to obtain a better representation of the true objective function as well (as seen in Fig. 11).
- 3) HRMS often realizes these improvements with fewer total function evaluations than standard sampling approaches (see Fig. 9). This is a characteristic that is critical in optimization of “expensive” objective functions (such as high-fidelity air combat simulations).

Having demonstrated the efficacy of RS and more generally HRMS on a highly noisy objective function, we next assess whether these insights translate to the adaptive training of AI pilots for the simulated air combat application.

B. Air Combat Simulation Experiments

The AI pilot here has many parameters that can be modified (see Table 1). However, only two parameters in particular were found to have significant effects on the outcome of the combat simulations for the one-on-one engagement scenario described in Sec. III: *launch* (launch delay) and *intspeed* (intercept speed). Hence, the simulation studies only focus on optimizing the two parameters $\mathbf{x}_b = [\text{launch}, \text{intspeed}]^T$ for the blue AI pilot to minimize the TTK metric, while fixing the others at their default values.

The simulation was optimized during 10 independent runs for each RS/MS GPBO configuration using 20 random initial seed points in the ranges $x_1 = (0, 5)$ and $x_2 = (0, 500)$ for each dimension of the search space. MAP hyperparameter estimation was used with the Laplace approximation, and the Matérn 3/2 kernel hyperpriors set were as follows: the mean as $\mathcal{N}(0, 100^2)$ and the covariance parameters as $\mathcal{U}(\log 1, \log 3)$.

Figure 12 shows the estimated locations of the optimum TTK as well as the ground truth optimum location estimated from a brute force grid search over the blue parameter space \mathbf{x}_b using 10,000 simulations (results for the UCB acquisition function over 10 optimization runs at each HRMS configuration are shown only; the remaining acquisition functions show similar trends and are omitted to reduce clutter). The estimates for \mathbf{x}_b and y grow tighter together, and closer to the ground truth, as both the RS and MS configuration numbers become greater than 1. The configurations marked by marked rectangles highlight that methods using solely SS, RS, and MS underperform the method that combines both RS and MS greater than 1. This finding is similar for the EI and TS functions as well and reflects the results found for the 3-D Ackley function. From this figure, we can conclude that there are HRMS configurations that yield more repeatable optimization results than SS and that neither RS nor MS alone is clearly better.

Figure 13 depicts some examples of the final GPs obtained during time limited optimization (again, allowing faster methods to use more evaluations/iterations) for three different HRMS configurations. The far left column is the ground truth GP model that is obtained by training with several thousands of samples over the input space. The key insight is that the RS3/MS3 strategy yielded a GP that better represents the ground truth. This confirms some of the observations from the previous study of the Ackley function. Also as in the previous study, Fig. 14 illustrates that

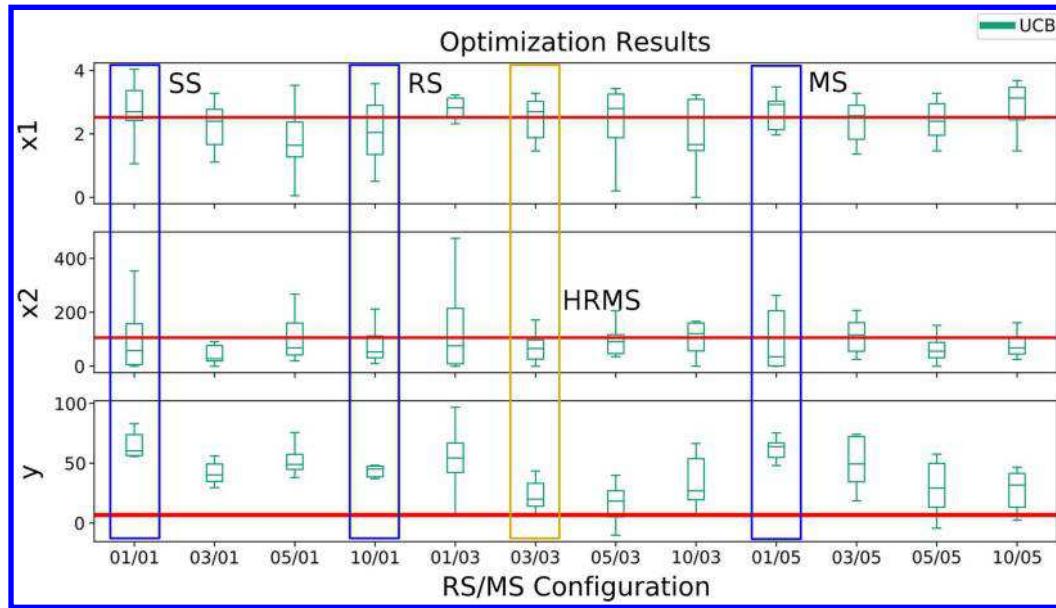


Fig. 12 Optimal x_1 , x_2 , and y = TTK for UCB vs HRMS configuration (solid horizontal line: true value).

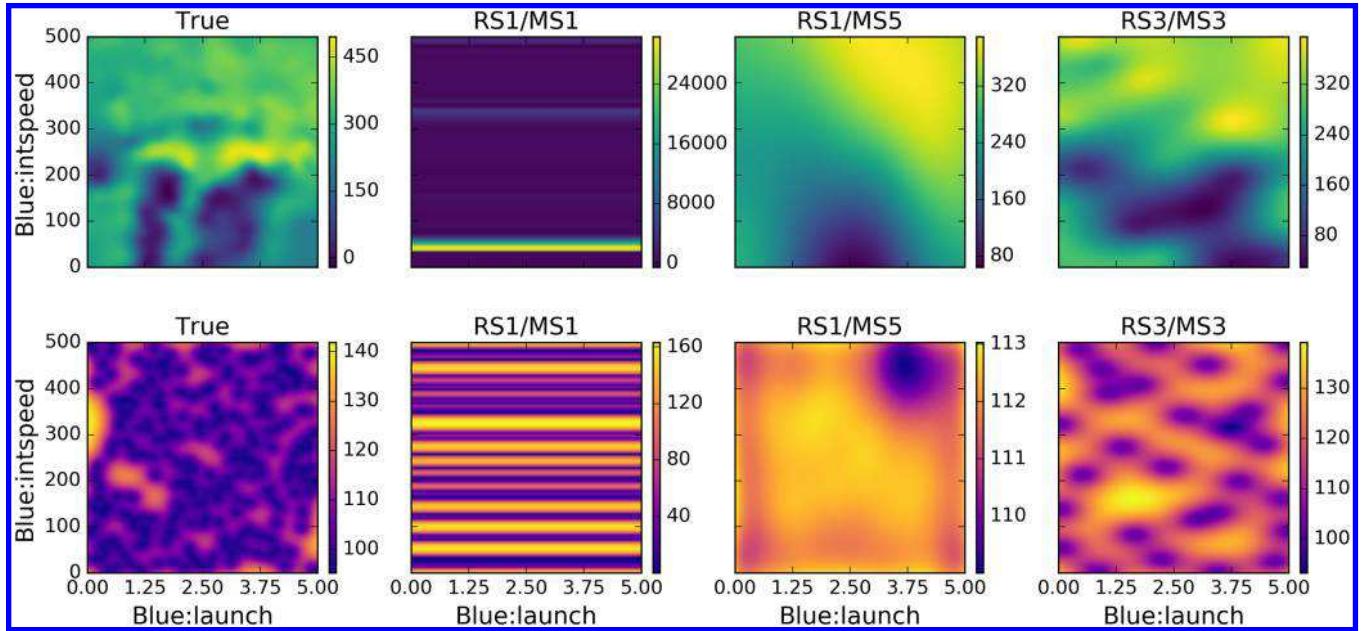


Fig. 13 Mean (top) and σ (bottom) for ‘true’ TTK and RS/MS GP estimates (after 1 h).

using HRMS improves the surrogate model as well as the optimization outcomes. These two phenomena are linked; a better surrogate model leads to more accurate and reliable optimization results. Returning to the high-level goal of this research, training an AI with behavioral parameters to optimize a combat objective and to provide tools for being adaptive, Fig. 12 demonstrates that GPBO is useful in this application. It shows that using HRMS results in more repeatable identification of the optimum AI behaviors in the simulated air combat.

Figure 13 shows that the combined RS3/MS3 strategy (column 4) yields a more accurate surrogate representation of the true objective function. This surrogate model gives insights about how the expected value of the TTK changes over the entire behavioral space for the blue AI agent. There seem to be two large areas with lower *intspeed* and moderate *launch* that provide the best performance for blue. Conversely, it appears that the higher *intspeed* configurations do not yield very good results. Looking at the full playback of recorded telemetry and decision data at the sampled simulation points in/near these regions helps gather more insight as to why.

Figure 15 shows the outcomes of two different simulations (rendered here in the open source TacView air combat simulation visualizer): a scenario with lower *launch* and *intspeed* values that has good results, and another scenario with higher *intspeed* that yields poor results because the blue agent cannot regain lock after launching a midrange missile. In this simulation, the aircraft must have an active lock on the target to continue guiding the missile. The images in the (A) column show the point at which the blue agent locks on and launches its midrange missile. The images in column (B) show that, in both scenarios, blue loses lock on red. However, in column (C), the blue with lower intercept speed is able to regain lock; on the bottom, blue is not able to maneuver to regain lock in time, and without lock, the missile is no longer guided to its target. As a result of blue missing red in the bottom scenario, it must spend a significant amount of time maneuvering again to regain lock on red. By lowering *launch* and *intspeed*, blue is able to maintain a lock on red and successfully guide its missile much more consistently on the first pass.

Although the RS3/MS3 strategy allowed some insight into the combat outcomes, the surrogate functions identified with SS and to a lesser extent RS1/MS5 (columns 2 and 3) offer less reliable insight into the true TTK function. The improved capability of predicting what might happen

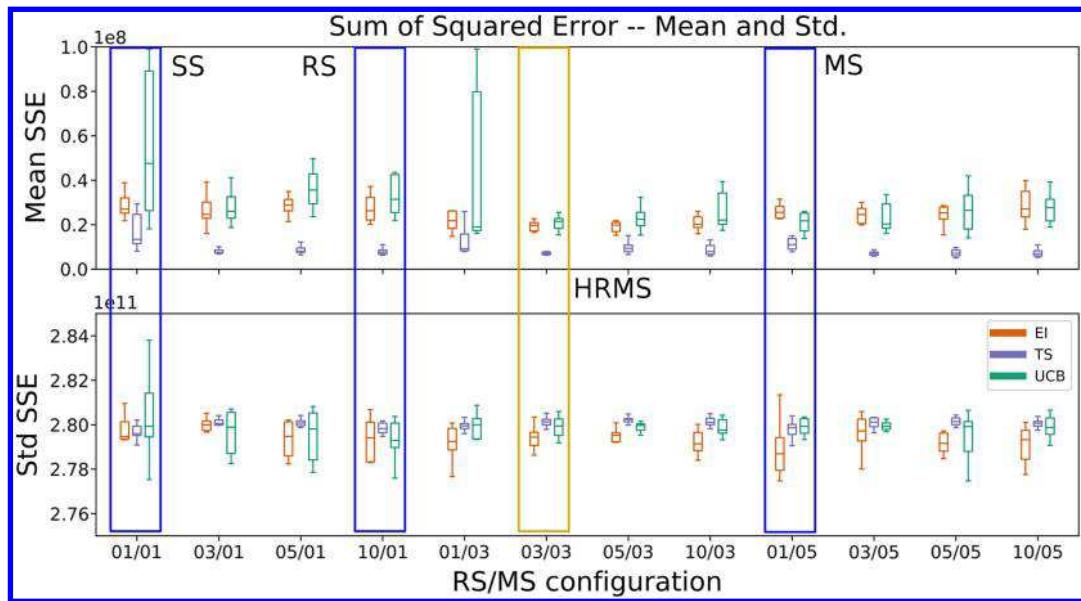


Fig. 14 SSE for TTK GP mean (top) and σ for different HRMS configurations.

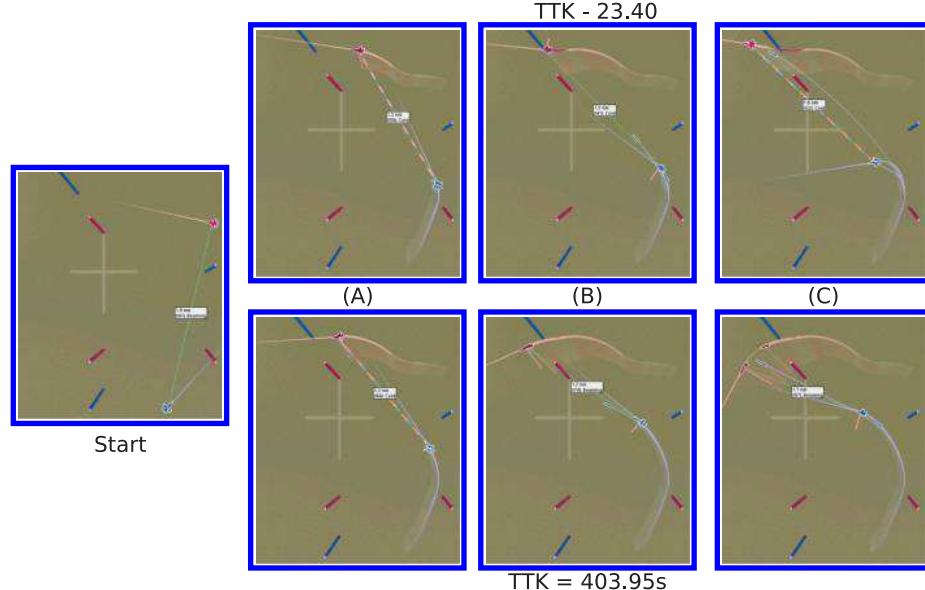


Fig. 15 Engagement outcomes (same start condition shown far left): $x_b = \{2.60, 125.0\}$, $y = \text{TTK} = 23.4$ s (top); $x_b = \{3.00, 450.0\}$, $y = \text{TTK} = 403.95$ s (bottom).

in untested conditions allows the AI to identify the expected optimum behavior but also to have more capacity to intelligently adapt by being able to predict outcomes at any other location in the behavioral space. Another ramification is that these models are, to some extent, interpretable by humans. It is easily observed from the two-dimensional plots in this case but might also be interpretable in higher dimensions by using more advanced visualization techniques.

C. Discussion

The simulations results show that GPBO with HRMS reliably finds optimum behaviors for the adaptive AI pilot, even given the inherent volatility in aerial combat. Compared to the pure RS and pure MS sampling conditions, HRMS-based GPBO is generally able to learn a more accurate model of the true objective function with fewer total combat simulations using any of the acquisition functions considered. The resulting model thus provides good predictions of engagement performance over the possible behavior/AI parameter solution space, thus allowing the AI to efficiently adapt behavior.

These results have several implications for minimizing the cost and improving the efficiency of mixed human–AI team LVC training. As already shown for AI-versus-AI training, parallelized simulation runs can be easily exploited to accelerate adaptation and behavior exploration via RS, MS, or HRMS. Although not explicitly demonstrated here, recent work in the computational cognitive science literature suggests that GPBO techniques are also quite effective tools for automating intensive training programs focused on evaluating and improving human performance in skills-based tasks [51]. As such, GPBO methods can be readily adapted to enhance human-versus-AI training as well. For instance, the blue force AI in the air combat mission simulation can be replaced by a human pilot, so that the adaptive AI now is responsible for the red agent instead. The goal of the GPBO learning engine in this case is then to maximize its performance (TTK or other metrics) against the human pilot with as few simulations as possible. The GP surrogate model produced by the learning engine would thus also serve as a valuable tool for predicting the human pilot's performance against

various red AI configurations. The learning engine could also be augmented to incorporate variations in combat environment parameters, so that the x search space for GPBO includes factors such as weather, wind, terrain, visibility, and others in addition to red agent parameters. For single human pilot training, the RS, MS, and HRMS methods can be used to repeatedly generate block test conditions, whose elements are then evaluated sequentially with the same pilot (possibly in randomized order). To fully leverage the parallelizability of GPBO sampling and exploration, these sampling techniques could also be used to generate an “average human” performance model using the outcomes of parallel engagements with multiple human pilots (e.g., who simultaneously engage the AI one-on-one in different instances of the same mission). It is interesting to note that the stochastic GPBO surrogate model can theoretically accommodate the high performance variability likely to be encountered among human pilots (due to individual differences, learning effects, etc.). An interesting open technical question in this case is the extent to which GPBO needs to be modified (if at all) to cope with possible nonstationarities in the objective function (i.e., changing performance statistics with respect to time and/or x parameters).

VI. Conclusions

This paper studied and further developed Gaussian process Bayesian optimization (GPBO) techniques for automatically tuning the decision-making parameters of an AI-controlled dogfighting agent, so that it could learn to optimize performance against an intelligent adversary. Performance was measured by a non-closed-form stochastic objective function evaluated on full-length simulated combat engagements and was modeled over the full agent behavior parameter space by a Gaussian process (GP) surrogate function. As per the standard GPBO formulation, the GP surrogate function was used to investigate new simulations at different parameter settings that most likely yield optimum performance. However, because of the highly volatile nature of the underlying objective function in this application, standard GPBO methods tend not to construct accurate surrogate models during optimization and thus tend not to yield reliable optimization results. Novel repeat sampling (RS) and hybrid repeat/multipoint sampling (HRMS) methods for GPBO were developed to address these issues. Numerical simulations on a benchmark optimization problem and a one-on-one AI-versus-AI dogfighting application scenario showed that GPBO with HRMS generally led to more accurate surrogate models and more efficient adaptation to the optimum agent parameters.

GPBO provides an attractive solution for performing optimization on expensive black-box objective functions, due to its ability to make predictions about the value of the objective function in unexplored locations. In many respects, the adaptive nature of GPBO closely resembles the explore-exploit behavior of reinforcement learning [52], with one major difference being that GPBO is applied here in between full engagements to optimize a holistic cost function in an “offline” sense, whereas reinforcement learning strategies might typically be applied during engagements to maximize total expected discounted reward functions (agent utilities) in an “online” sense (i.e., updating agent behaviors in the middle of an engagement). Although the connections between reinforcement learning and GPBO have been noted elsewhere [24], the concept and motivations behind using RS and HRMS formulations of GPBO have not been considered previously. The results provided here show a clear benefit for improving the quality of the surrogate model and reliability of the parameter search process using a minimum number of function evaluations (combat simulations). However, although this work showed some analysis of the effects of RS compared to traditional single sampling and multipoint sampling (MS) GPBO methods, additional theoretical work is needed to fully understand the benefits of RS and HRMS strategies. This deeper understanding can be used to develop techniques for automatically calculating RS and MS sample size portions in HRMS, to ensure best possible optimization performance without guesswork or brute force evaluations to determine the best configurations. The existing literature regarding statistical experiment design may be helpful to understanding the contributions of RS in this regard.

Although this work only focused on simulated engagements between two AIs, the GPBO methods developed here can, in principle, be extended to simulations with human participants. This work suggests many interesting pathways for exploring how GPBO techniques can be used to generate adaptive red agents for human trainees. In addition, the GPBO methods developed here could be used to develop “backseat coaches” for human blue force operators (i.e., AI decision aids that suggest behaviors that will most likely yield successful results based on observed performance). The problem of simultaneously training multiple red and blue force agents presents many other interesting possibilities for extending the GPBO techniques developed here. To cope with the high dimensionality and sensitivity to hyperparameter settings in such applications, it may be necessary to combine RS and HRMS approaches with other sophisticated GP regression and acquisition function optimization strategies not considered here, such as fully Bayesian hyperparameter learning [31] and “random embedding” GPBO for very high-dimensional problems [40].

References

- [1] Doyle, M. J., and Portrey, A. M., “Rapid Adaptive Realistic Behavior Modeling is Viable for Use in Training,” *Proceedings of the 23rd Conference on Behavior Representation in Modeling and Simulation*, BRiMS Committee, Curran Associates, Inc., 2014, pp. 73–80.
- [2] McManus, J. W., and Goodrich, K. H., “Artificial Intelligence (AI) Based Tactical Guidance for Fighter Aircraft,” NASA TR Rept. AIAA-90-3435, 1990.
- [3] Moore, S. B., Madison, W. G., Sepp, G. D., Stracener, J. T., and Coward, R. E., “Air Combat Training: Good Stick Index Validation,” Vought Corp., TR AFHRL-TR-79-15, Dallas, TX, June 1979.
- [4] Kelly, M. J., “Performance Measurement During Simulated Air-to-Air Combat,” *Human Factors: The Journal of the Human Factors and Ergonomics Society*, Vol. 30, No. 4, 1988, pp. 495–506.
doi:10.1177/00187208803000410
- [5] Mulgund, S., Harper, K., Krishnakumar, K., and Zacharias, G., “Air Combat Tactics Optimization Using Stochastic Genetic Algorithms,” *Proceedings of 1998 IEEE International Conference on Systems, Man, and Cybernetics*, Vol. 4, IEEE Publ., Piscataway, NJ, Oct. 1998, pp. 3136–3141.
- [6] Huynh, H. T., and Costes, P., *Numerical Optimization of Air Combat Maneuvers*, AIAA, New York, Aug. 1987, pp. 647–658.
- [7] Mulgund, S., Harper, K., and Zacharias, G., “Large-Scale Air Combat Tactics Optimization Using Genetic Algorithms,” *Journal of Guidance, Control, and Dynamics*, Vol. 24, No. 1, Jan. 2001, pp. 140–142.
doi:10.2514/2.4689
- [8] Paranjape, A. A., and Ananthkrishnan, N., “Combat Aircraft Agility Metrics—A Review,” *Journal of Aerospace Sciences and Technologies*, Vol. 58, No. 2, 2006, pp. 143–154.
- [9] Luo, D.-L., Shen, C.-L., Wang, B., and Wu, W.-H., “Air Combat Decision-Making for Cooperative Multiple Target Attack Using Heuristic Adaptive Genetic Algorithm,” *Proceedings of the 4th International Conference on Machine Learning and Cybernetics*, Vol. 1, Springer-Verlag, Berlin, Aug. 2005, pp. 473–478.
- [10] Gonsalves, P., and Burge, J., “Software Toolkit for Optimizing Mission Plans (STOMP),” *AIAA 1st Intelligent Systems Technical Conference*, AIAA Paper 2004-6279, Sept. 2004.
- [11] Ernest, N., Carroll, D., Schumacher, C., Clark, M., and Cohen, K., “Genetic Fuzzy Based Artificial Intelligence for Unmanned Combat Aerial Vehicle Control in Simulated Air Combat Missions,” *Journal of Defense Management*, Vol. 6, No. 1, 2016, pp. 1–3.
- [12] Yannakakis, G. N., “Game AI Revisited,” *Proceedings of the 9th Conference on Computing Frontiers*, ACM Press, New York, 2012, p. 285.
- [13] Cole, N., Louis, S. J., and Miles, C., “Using a Genetic Algorithm to Tune First-Person Shooter Bots,” *Proceedings of the 2004 Congress on Evolutionary Computation*, Vol. 1, IEEE Publ., Piscataway, NJ, June 2004, pp. 139–145.
- [14] Liaw, C., Wang, W.-H., Tsai, C.-T., Ko, C.-H., and Hao, G., “Evolving a Team in a First-Person Shooter Game by Using a Genetic Algorithm,” *Applied Artificial Intelligence*, Vol. 27, No. 3, March 2013, pp. 199–212.
doi:10.1080/08839514.2013.768883

[15] Othman, N., Decraene, J., Cai, W., Hu, N., Low, M. Y. H., and Gouaillard, A., "Simulation-Based Optimization of StarCraft Tactical AI Through Evolutionary Computation," *Proceedings of the 2012 IEEE Conference on Computational Intelligence and Games*, IEEE Publ., Piscataway, NJ, Sept. 2012, pp. 394–401.

[16] Kelly, M. J., Wooldridge, L., Hennessy, R. T., Vreuls, D., Barnebey, S. F., Cotton, J. C., and Reed, J. C., "Air Combat Maneuvering Performance Measurement," Canyon Research Group TR 1, Westlake Village, CA, Sept. 1979.

[17] Rasmussen, C. E., and Williams, C. K. I., *Gaussian Processes for Machine Learning*, Adaptive Computation and Machine Learning, MIT Press, Cambridge, MA, 2006, pp. 1–272.

[18] Hutter, F., Hoos, H. H., and Leyton-Brown, K., "Sequential Model-Based Optimization for General Algorithm Configuration," *Lecture Notes in Computer Science (Including Subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*, Vol. 6683, Springer, Berlin, 2011, pp. 507–523.

[19] Snoek, J., Rippel, O., Swersky, K., Kiros, R., Satish, N., Sundaram, N., Patwary, M. M. A., Prabhat, , and Adams, R. P., "Scalable Bayesian Optimization Using Deep Neural Networks," *Proceedings of the 32nd International Conference on Machine Learning*, Feb. 2015, pp. 2171–2180, PMLR.org.

[20] Gelbart, M. A., "Constrained Bayesian Optimization and Applications," Ph.D. Thesis, Harvard Univ., Cambridge, MA, May 2015.

[21] Eggensperger, K., Feurer, M., Hutter, F., Bergstra, J., Snoek, J., Hoos, H. H., and Leyton-Brown, K., "Towards an Empirical Foundation for Assessing Bayesian Optimization of Hyperparameters," *NIPS Workshop on Bayesian Optimization in Theory and Practice*, 2013, p. 5, <https://bayesopt.github.io/past/2013.html>.

[22] Bishop, C. C. M., *Pattern Recognition and Machine Learning*, Vol. 4, Information Science and Statistics, Springer, New York, 2006, pp. 1–738.

[23] Shahriari, B., Swersky, K., Wang, Z., Adams, R. P., and de Freitas, N., "Taking the Human Out of the Loop: A Review of Bayesian Optimization," *Proceedings of the IEEE*, Vol. 104, No. 1, 2016, pp. 148–175.
doi:10.1109/JPROC.2015.2494218

[24] Brochu, E., Cora, V. M., and De Freitas, N. L. B. B., "A Tutorial on Bayesian Optimization of Expensive Cost Functions, with Application to Active User Modeling and Hierarchical Reinforcement Learning," arXiv preprint arXiv:1012.2599, 2010, <https://arxiv.org/abs/1012.2599>.

[25] Hoffman, M., Brochu, E., and Freitas, N. D., "Portfolio Allocation for Bayesian Optimization," *Proceedings of the Conference on Uncertainty in Artificial Intelligence*, Morgan Kaufmann Publishers Inc., San Francisco, CA, 2011, pp. 327–336.

[26] Jones, D. R., Schonlau, M., and William, J., "Efficient Global Optimization of Expensive Black-Box Functions," *Journal of Global Optimization*, Vol. 13, No. 4, 1998, pp. 455–492.
doi:10.1023/A:1008306431147

[27] Srinivas, N., Krause, A., Kakade, S. M., and Seeger, M., "Gaussian Process Optimization in the Bandit Setting: No Regret and Experimental Design," arXiv preprint arXiv: 0912.3995, June 2010.
doi:10.1109/TIT.2011.2182033

[28] Johnson, S. G., "The NLOpt Nonlinear-Optimization Package," GitHub, May 2014, <http://ab-initio.mit.edu/nlopt> [accessed 01 Oct. 2015].

[29] Forrester, A., Sobester, A., and Keane, A., *Engineering Design via Surrogate Modelling: A Practical Guide*, Wiley, Chichester, U.K., 2008, pp. 1–228.

[30] Bergstra, J., Bardenet, R., Bengio, Y., and Kegl, B., "Algorithms for Hyper-Parameter Optimization," *Proceedings of the 24th International Conference on Neural Information Processing Systems, NIPS'11*, Curran Associates Inc., Red Hook, NY, 2011, pp. 2546–2554, <http://dl.acm.org/citation.cfm?id=2986459.2986743>.

[31] Snoek, J., Larochelle, H., and Adams, R. P., "Practical Bayesian Optimization of Machine Learning Algorithms," *Proceedings of the 25th International Conference on Neural Information Processing Systems, NIPS'12*, Curran Associates Inc., Red Hook, NY, 2012, pp. 2951–2959.

[32] Mahendran, N., Wang, Z., Hamze, F., and De Freitas, N., "Adaptive MCMC with Bayesian Optimization," *Proceedings of the International Conference on Artificial Intelligence and Statistics*, Vol. 9, 2012, p. 152, PMLR.org.

[33] Bergstra, J., and Bengio, Y., "Random Search for Hyper-Parameter Optimization," *Journal of Machine Learning Research*, Vol. 13, No. 1, 2012, pp. 281–305.

[34] Macdonald, I. A., "Comparison of Sampling Techniques on the Performance of Monte-Carlo Based Sensitivity Analysis," *Proceedings of the 11th Conference of International Building Performance Simulation Association*, International Building Performance Simulation Association, Curran Associates, Red Hook, NY, July 2009, pp. 992–999.

[35] Schonlau, M., "Computer Experiments and Global Optimization," Ph.D. Thesis, Univ. of Waterloo, Waterloo, ON, Canada, 1997.

[36] Hoffman, M. W., and Shahriari, R., "Modular Mechanisms for Bayesian Optimization," *Proceedings of the NIPS Workshop on Bayesian Optimization*, 2014, pp. 1–5.

[37] Chevalier, C., and Ginsbourger, D., "Fast Computation of the Multi-Points Expected Improvement with Applications in Batch Selection," *Lecture Notes in Computer Science (Including Subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*, Vol. 7997, Springer, Berlin, 2013, pp. 59–69.

[38] Huang, D., "Experimental Planning and Sequential Kriging Optimization Using Variable Fidelity Data," Ph.D. Thesis, Ohio State Univ., Columbus, OH, 2005.

[39] Ponweiser, W., Wagner, T., and Vincze, M., "Clustered Multiple Generalized Expected Improvement: A Novel Infill Sampling Criterion for Surrogate Models," *Proceedings of the 2008 IEEE Congress on Evolutionary Computation*, IEEE Publ., Piscataway, NJ, June 2008, pp. 3515–3522.

[40] Wang, Z., Zoghi, M., Hutter, F., Matheson, D., and Freitas, N. D., "Bayesian Optimization in High Dimensions via Random Embeddings," *Proceedings of the 23rd International Joint Conference on Artificial Intelligence*, AAAI Press, Palo Alto, CA, 2013, pp. 1778–1784.

[41] Jones, D. R., Law, C., and Law, C., "Lipschitzian Optimization Without the Lipschitz Constant," *Journal of Optimization Theory and Applications*, Vol. 79, No. 1, 1993, pp. 157–181.
doi:10.1007/BF00941892

[42] Cawley, G. C., and Talbot, N. L. C., "Preventing Over-Fitting During Model Selection via Bayesian Regularisation of the Hyper-Parameters," *Journal of Machine Learning Research*, Vol. 8, April 2007, pp. 841–861.

[43] Cawley, G. C., and Talbot, N. L. C., "On Over-Fitting in Model Selection and Subsequent Selection Bias in Performance Evaluation," *Journal of Machine Learning Research*, Vol. 11, Aug. 2010, pp. 2079–2107.

[44] Contal, E., Buffoni, D., Robicquet, A., and Vayatis, N., "Parallel Gaussian Process Optimization with Upper Confidence Bound and Pure Exploration," *Lecture Notes in Computer Science (Including Subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*, Vol. 8188, Springer, Berlin, 2013, pp. 225–240.

[45] Ginsbourger, D., Riche, R. L., and Carraro, L., "Kriging Is Well-Suited to Parallelize Optimization," *Computational Intelligence in Expensive Optimization Problems*, Vol. 2, No. 2, 2010, pp. 131–162.
doi:10.1007/978-3-642-10701-6

[46] Pyzdek, T., and Keller, P. A., *Quality Engineering Handbook*, CRC Press, Boca Raton, FL, 2003, pp. 1–732.

[47] Thompson, W. R., "On the Likelihood That One Unknown Probability Exceeds Another in View of the Evidence of Two Samples," *Biometrika*, Vol. 25, Nos. 3–4, Dec. 1933, p. 285.

[48] Wang, J., Clark, S. C., Liu, E., and Frazier, P. I., "Parallel Bayesian Global Optimization of Expensive Functions," arXiv preprint arXiv:1602.05149, Feb. 2016, <https://arxiv.org/abs/1602.05149>.

[49] Rasmussen, C. E., and Williams, C. K. I., "GPML Toolbox," ebrary, Dec. 2006, <http://www.gaussianprocess.org/gpml/code/matlab/doc/> [retrieved 01 Sept. 2015].

[50] Back, T., *Evolutionary Algorithms in Theory and Practice: Evolution Strategies, Evolutionary Programming, Genetic Algorithms*, Oxford Univ. Press, New York, Jan. 1996, pp. 1–328.

[51] Khajah, M. M., Roads, B. D., Lindsey, R. V., Liu, Y.-E., and Mozer, M. C., "Designing Engaging Games Using Bayesian Optimization," *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems*, ACM, New York, 2016, pp. 5571–5582.

[52] Russell, S., and Norvig, P., *Artificial Intelligence: A Modern Approach*, Prentice-Hall, Upper Saddle River, NJ, 2003, pp. 1–1152.