

Gaussian Max-Value Entropy Search for Multi-Agent Bayesian Optimization

Haitong Ma, Tianpeng Zhang, Yixuan Wu, Flavio P. Calmon, Na Li

Abstract—We study the multi-agent Bayesian optimization (BO) problem, where multiple agents maximize a black-box function via iterative queries. We focus on Entropy Search (ES), a sample-efficient BO algorithm that selects queries to maximize the mutual information about the maximum of the black-box function. One of the main challenges of ES is that calculating the mutual information requires computationally-costly approximation techniques. For multi-agent BO problems, the computational cost of ES is exponential in the number of agents. To address this challenge, we propose the *Gaussian Max-value Entropy Search*, a multi-agent BO algorithm with favorable sample and computational efficiency. The key to our idea is to use a normal distribution to approximate the function maximum and calculate its mutual information accordingly. The resulting approximation allows queries to be cast as the solution of a closed-form optimization problem which, in turn, can be solved via a modified gradient ascent algorithm and scaled to a large number of agents. We demonstrate the effectiveness of Gaussian max-value Entropy Search through numerical experiments on standard test functions and real-robot experiments on the source seeking problem. Results show that the proposed algorithm outperforms the multi-agent BO baselines in the numerical experiments and can stably seek the source with a limited number of noisy observations on real robots.

I. INTRODUCTION

Bayesian optimization (BO) is a sample-efficient method for maximizing expensive-to-evaluate black-box functions, which frequently arise in robotic applications. By iteratively evaluating the black-box function at the *query points*, BO first builds a probabilistic model about the function and then infers the location of the maximum accordingly. Scenarios where BO is applicable include tuning the parameters of controllers [1] and motion planners [2], [3], seeking the location of source signal [4], [5], designing the morphology structure of robots [6], and so on.

Query point selection is a fundamental challenge in BO. The key is to carefully choose the query points to learn about the objective function (exploration) while leveraging existing knowledge to maximize it (exploitation). One approach of BO is to select query points in each iterative step to maximize the so-called *acquisition function*. The acquisition function regulates the tension between exploring versus exploiting and is updated at each iteration based on all the queries and observations collected so far. One of the most notable acquisition functions is the Gaussian process upper confidence bound (GP-UCB) [7], [8] extended from the multi-armed bandit problem. GP-UCB uses a weighted sum of

the posterior mean (exploitation) and the posterior variance (exploration) of the Gaussian process model to handle the trade-off. Other popular acquisition functions include the probability of improvement (PI) [9], expected improvement (EI) [10], knowledge gradient (KG) [11], Entropy Search (ES) [12], Thompson sampling (TS) [13], etc.

While most BO methods use a single agent for querying, it is more desirable to have multiple agents querying the black-box function simultaneously in many applications. With multiple agents, the querying is parallelized so that more information can be obtained per iteration, and thus the objective function can be learned and maximized faster. Existing multi-agent BO studies have proposed two approaches to selecting the batch of query points for the agents at each iterative step: (1) sequential query calculation and (2) batch query calculation. Sequential query calculation computes the query of each agent within the batch one by one. A single-agent BO algorithm usually determines the first agent's query. Then, other agents' queries are added sequentially to provide more exploration [14]–[16] or exploitation [17], [18]. In contrast, batch query calculation computes the batch of queries simultaneously, for example, using Thompson sampling or Entropy Search [19]–[21]. Apart from the distinctions in the computational procedure, these two approaches also address the collaboration among the agents differently. The collaboration can be viewed as balancing the exploration-exploitation trade-off within the same batch of query points. While the sequential query calculation usually assigns explicit roles of exploration/exploitation to each agent, the batch query calculation handles the collaboration implicitly through its probabilistic model.

Among these multi-agent methods, the Entropy Search (ES) has gained increasing attention because of its promising low-regret performance [12], [22]–[25]. ES maximizes the black-box function by maximizing the *mutual information* of the estimated function maximum, which is shown to be more sample-efficient than directly querying the function at the estimated function maximum [12]. Moreover, ES is especially suitable for the multi-agent BO setting since the collaboration among the agents can be encouraged by maximizing the total mutual information about the function maximum in the agents' queries. In this way, ES can automatically adapt the multi-agent collaboration under different objective functions to obtain the most informative queries, which is more flexible and efficient than the fixed role assignment scheme used in the sequential query calculation methods.

One of the main challenges of ES, especially for multi-agent BO problems, is computational efficiency. Computing the entropy of the function maximum is generally intractable and requires sophisticated approximation techniques. Exist-

The authors are with School of Engineering and Applied Sciences, Harvard University. Email: {haitongma, tzhang, florencewu, fcalmon}@g.harvard.edu, nali@seas.harvard.edu. The work is supported under ONR YIP: N00014-19-1-2217, NSF CNS: 2003111, and NSF AI institute: 2112085.

ing methods, including Monte-Carlo sampling [22], expectation propagation [12], [20], [25], random feature sampling [20], [23], [24], and Gambel sampling [23], have exponential computational cost in the number of agents [20], [25]. Although [20] proposes a gradient-based multi-agent ES algorithm to reduce the exponential cost to polynomial cost, the computation is still heavy since the gradient calculation involves a large number of matrix inversions. The required number of inversions is proportional to the dataset size.

Our Contributions. We proposed the Gaussian Max-value Entropy Search (GMES), a computationally efficient multi-agent Entropy Search algorithm with a novel entropy approximation scheme and practical implementations for the multi-agent setting. Specifically, we use the normal distribution to approximate the distribution of the function maximum and calculate its entropy. We use the mutual information for this approximate distribution as our acquisition function, which has a closed-form expression. Unlike existing multi-agent Entropy Search algorithms, we do not need costly sampling when calculating the acquisition function. We further use gradient ascent to compute the query points and add log-barrier safety constraints to make the proposed algorithm scalable to a large number of agents and applicable to real-world applications. Together, our algorithm has favorable computational efficiency compared to existing methods.

We then test the algorithms on both numerical and real-robot experiments. Experiment results show that the proposed algorithm outperforms the baseline multi-agent BO algorithms with different numbers of agents in numerical tests. The real-robot experiments demonstrate that our algorithm can successfully seek a light source with a small number of queries. These source seeking experiments also showcase the substantial advantage of using multiple agents over a single agent. Compared to single-agent seeking, four agents improved the source-seeking time by 59.9% and the source-seeking iterations by 67.6% on average.

II. PRELIMINARIES

A. Problem Formulation

We consider a multi-agent BO problem that consists of a team of m agents, a compact domain $\mathcal{X} \subset \mathbb{R}^d$, and an unknown function

$$f: \mathcal{X} \rightarrow \mathbb{R}.$$

We assume f is continuous on $\mathcal{X}$, so its maximum exists in this domain. The goal is to maximize f only through the queries of function values. We assume each agent can query f at any point $\mathbf{x} \in \mathcal{X}$ and observe a noisy function value

$$\mathbf{y} = f(\mathbf{x}) + \epsilon, \quad \epsilon \sim \mathcal{N}(0, \sigma_0^2),$$

where σ_0^2 is the variance of observation noise ϵ .

The agents query f through a sequence of iterations $t = 1, 2, \dots, T$. For all agents $i \in \{1, 2, \dots, m\}$, we use lower-case letters $\mathbf{x}_t^i$ for the query point of agent i at time t and $\mathbf{y}_t^i$ for agent i 's observation at time t . We use capital letters $X_t \equiv \{\mathbf{x}_t^1, \mathbf{x}_t^2, \dots, \mathbf{x}_t^m\}$ and $Y_t \equiv \{\mathbf{y}_t^1, \mathbf{y}_t^2, \dots, \mathbf{y}_t^m\}$ to denote the collections of agents' query points and evaluations at time t . We denote $\mathcal{X}^m = \underbrace{\mathcal{X} \times \mathcal{X} \times \dots \times \mathcal{X}}_m \in \mathbb{R}^{md}$ as the domain

of batch queries. We define $\mathbb{X}_t \equiv \{X_1, X_2, \dots, X_t\}$ and $\mathbb{Y}_t \equiv \{Y_1, Y_2, \dots, Y_t\}$ as all the queries and observations up to time t . Let $D_t = \mathbb{X}_{t-1} \cup \mathbb{Y}_{t-1}$ be the observed data before time t . We assume the agents know and only know D_t when deciding the query points X_t .

B. Gaussian Process

We briefly introduce the Gaussian process (GP), our probabilistic model of the objective function. A GP model is built upon the observed data D_t and a positive-definite kernel function $k(\mathbf{x}, \mathbf{x}')$ that models our prior belief about the coupling between $f(\mathbf{x})$ and $f(\mathbf{x}')$ [26]. Given D_t and k , the GP model for function f can be fully described by its mean value function $\mu_t(\mathbf{x})$ and covariance function $\Sigma_t(\mathbf{x}, \mathbf{x}')$, which are calculated by

$$\begin{aligned} \mu_t(\mathbf{x}) &= \mathbf{k}_t(\mathbf{x})^\top (\mathbf{K}_t + \sigma_0^2 \mathbf{I})^{-1} \mathbf{y}_t, \\ \Sigma_t(\mathbf{x}, \mathbf{x}') &= k(\mathbf{x}, \mathbf{x}') - \mathbf{k}_t(\mathbf{x})^\top (\mathbf{K}_t + \sigma_0^2 \mathbf{I})^{-1} \mathbf{k}_t(\mathbf{x}') \end{aligned} \quad (1)$$

where $\mathbf{k}_t(\mathbf{x}) = [k(\mathbf{x}', \mathbf{x})]_{\mathbf{x}' \in \mathbb{X}_{t-1}}^\top$ and $\mathbf{y}_t = [\mathbf{y}]_{\mathbf{y} \in \mathbb{Y}_{t-1}}^\top$ are $m \cdot (t-1)$ dimensional vectors, and $\mathbf{K}_t = [k(\mathbf{x}, \mathbf{x}')]_{\mathbf{x}, \mathbf{x}' \in \mathbb{X}_{t-1}}$ is a $m \cdot (t-1) \times m \cdot (t-1)$ positive definite matrix. $\mathbf{I}$ is the identity matrix with the same shape as $\mathbf{K}_t$. We denote the resulting GP model as $GP(\mu_t, \Sigma_t | D_t)$. We also define $\sigma_t^2(\mathbf{x}) = \Sigma_t(\mathbf{x}, \mathbf{x})$ as the variance function, and let $\sigma_t(\mathbf{x}) = \sqrt{\sigma_t^2(\mathbf{x})}$. The value $\mu_t(\mathbf{x})$ can be understood as the GP model's predicted value of $f(\mathbf{x})$ given the observed data D_t and $\sigma_t^2(\mathbf{x})$ is the uncertainty in that prediction in the form of variance. A more comprehensive discussion of topics related to GP can be found in [26].

C. Posterior Max-value and Entropy Search

From the function space viewpoint, GPs can be regarded as distributions over functions [26]. Denote the random function sampled from posterior GP distribution conditioned on the observed data D_t by $\hat{f}_t | D_t \sim GP(\mu_t, \Sigma_t | D_t)$. The corresponding maximal value of $\hat{f}_t$ is thus a random variable, which we denote as

$$\mathbf{y}^{\max} | D_t \sim \max_{\mathbf{x} \in \mathcal{X}} \hat{f}_t(\mathbf{x}) | D_t \quad (2)$$

We name $\mathbf{y}^{\max} | D_t$ as the *posterior max-value*. The key of Entropy Search (ES) is to select query points that maximally reduce the uncertainty in $\mathbf{y}^{\max}$, and specifically, the uncertainty is quantified by the differential entropy of $\mathbf{y}^{\max}$, defined by

$$H(\mathbf{y}^{\max} | D_t) = - \int p(\mathbf{y}^{\max} | D_t) \log(p(\mathbf{y}^{\max} | D_t)) \, d\mathbf{y}^{\max}.$$

The uncertainty about $\mathbf{y}^{\max}$ decreases after gaining information from the new queries and observations $\{X_t, Y_t\}$. The reduction in uncertainty, or mutual information, is defined by the entropy reduction after conditioning on more information,

$$I(\mathbf{y}^{\max}; \{X_t, Y_t\} | D_t) \quad (3)$$

$$= H(\mathbf{y}^{\max} | D_t) - H(\mathbf{y}^{\max} | D_t \cup \{X_t, Y_t\}) \quad (4)$$

$$= H(X_t, Y_t | D_t) - H(X_t, Y_t | D_t, \mathbf{y}^{\max}) \quad (5)$$

where $I(\mathbf{y}^{\max}; \{X_t, Y_t\} | D_t)$ denotes the mutual information between $\mathbf{y}^{\max}$ and new data $\{X_t, Y_t\}$ to be added.

Equation (4) follows from the definition of mutual information and Equation (5) is from the symmetric property of mutual information [27]. In summary, ES uses the mutual information above as the acquisition function to quantify the uncertainty reduction about $\mathbf{y}^{\max}$. The agents query X_t that maximizes (3) to get the most information about $\mathbf{y}^{\max}$.

The key challenge in ES is *how to calculate the mutual information* in (3) since it usually does not have a closed form. Equations (4) and (5) represent two different approaches, and both have difficulties. The difficulty in (4) is that the differential entropy $H(\mathbf{y}^{\max} | D_t \cup \{X_t, Y_t\})$ is conditioned on Y_t which is not known before querying the function at X_t [12], [22]. One solution proposed by [28] is to sample a batch of realizations $\{\hat{Y}_{1,t}, \hat{Y}_{2,t}, \dots\}$ from the posterior distribution $Y_t | D_t$ and compute $H(\mathbf{y}^{\max} | D_t \cup \{X_t, Y_t\})$ by averaging over $\mathbb{H}_t = \{H(\mathbf{y}^{\max} | D_t \cup \{X_t, Y_t = \hat{Y}_{i,t}\}) : i = 1, 2, \dots\}$. However, the corresponding computational cost is high since for each $\hat{Y}_{i,t}$ we need to calculate $H(\mathbf{y}^{\max} | D_t \cup \{X_t, Y_t = \hat{Y}_{i,t}\})$ once, and the calculation must be done using an iterative algorithm. Therefore, the total computation for (4) is substantial.

To avoid conditioning on Y_t , recent studies have considered computing the mutual information with equation (5) [20], [23]–[25]. However, this approach still requires sophisticated computation because the term $H(X_t, Y_t | D_t, \mathbf{y}^{\max})$ is non-trivial due to the conditioning on $\mathbf{y}^{\max}$.

The multi-agent nature of our problem brings further computation challenges to ES. Computing the mutual information through either Eq. (4) or Eq. (5) typically requires averaging over a sufficient number of samples from the m -dimensional posterior distribution $Y_t | D_t$. Therefore, the number of samples needed for a good approximation of the mutual information is exponential in the number of agents m . Furthermore, optimizing the non-trivial entropy over X_t on the $m \times d$ -dimensional space introduces even more difficulties [25]. Brute-force methods, such as grid or random search, also induce exponential optimization cost (in m).

In this paper, we follow the approach of (4) to calculate the acquisition function. We aim to reduce the computation cost and leverage the information-theoretic collaboration scheme to develop an efficient multi-agent Entropy Search algorithm that is practical for real robotic applications. We use the normal distribution to approximate the distribution of $\mathbf{y}^{\max}$ and use the mutual information associated with the approximated distribution as our acquisition function. We show that our acquisition function has an explicit expression; its computation is thus free from the expensive sampling over the m -dimensional distribution $Y_t | D_t$ in previous approaches like [28]. With the closed-form acquisition function, we develop a centralized multi-agent Entropy Search algorithm where multiple agents collaborate on a BO task. Gradient ascent is applied to mitigate the heavy optimization computation over X_t in the multi-agent setting.

III. GAUSSIAN MAX-VALUE ENTROPY SEARCH

A. Gaussian Approximation of Posterior Max-Value

Although the value $\hat{f}_t(\mathbf{x})$ follows normal distribution for each $\mathbf{x} \in \mathcal{X}$, the maximal value $\mathbf{y}^{\max} = \max_{\mathbf{x} \in \mathcal{X}} \hat{f}_t(\mathbf{x})$ is not Gaussian in general. Therefore, the differential entropies

of $\mathbf{y}^{\max}$ usually do not have a closed form. To overcome this difficulty, we use the distribution of a Gaussian random variable $\tilde{\mathbf{y}}^{\max}$ to approximate the posterior distribution $\mathbf{y}^{\max} | D_t$ so that the resulting mutual information associated with $\tilde{\mathbf{y}}^{\max}$ has a closed-form expression.

We formally define $\tilde{\mathbf{y}}^{\max}$ as follows. Denote the location with the maximal upper confidence bound(UCB) value at time t as

$$\mathbf{x}_t^{\text{ucb}} \equiv \operatorname{argmax}_{\mathbf{x} \in \mathcal{X}} (\mu_t(\mathbf{x}) + \beta_t \sigma_t(\mathbf{x}))$$

where β_t is a hyper-parameter indicating the confidence interval length. The UCB, $\mu_t(\mathbf{x}) + \beta_t \sigma_t(\mathbf{x})$, follows the definition in GP-UCB [8]. Define the *estimated* posterior max-value $\tilde{\mathbf{y}}^{\max} | D_t$ as

$$\begin{aligned} \tilde{\mathbf{y}}^{\max} | D_t &\sim \mathcal{N}(\mu_t(\mathbf{x}_t^{\text{ucb}}), \sigma_t^2(\mathbf{x}_t^{\text{ucb}})) \\ \tilde{\mathbf{y}}^{\max} | D_t \cup \{X_t, Y_t\} &\sim \mathcal{N}(\mu_{t+1}(\mathbf{x}_t^{\text{ucb}}), \sigma_{t+1}^2(\mathbf{x}_t^{\text{ucb}})) \end{aligned} \quad (6)$$

Here, $\mu_t(\mathbf{x}_t^{\text{ucb}})$, $\sigma_t^2(\mathbf{x}_t^{\text{ucb}})$ and $\mu_{t+1}(\mathbf{x}_t^{\text{ucb}})$ and $\sigma_{t+1}^2(\mathbf{x}_t^{\text{ucb}})$ are the t and $t+1$ step posterior mean and variance functions at $\mathbf{x}_t^{\text{ucb}}$. $\mu_t(\mathbf{x}_t^{\text{ucb}})$, $\sigma_t^2(\mathbf{x}_t^{\text{ucb}})$ are computed before we get the observation Y_t at X_t , and $\mu_{t+1}(\mathbf{x}_t^{\text{ucb}})$, $\sigma_{t+1}^2(\mathbf{x}_t^{\text{ucb}})$ are computed after we get the observation Y_t at X_t . This way, we can compare the entropy before and after observing Y_t to calculate the mutual information.

We approximate the mutual information in (3) with $I(\tilde{\mathbf{y}}^{\max}; \{X_t, Y_t\} | D_t)$, which is given by

$$I(\tilde{\mathbf{y}}^{\max}; \{X_t, Y_t\} | D_t) \quad (7)$$

$$= H(\tilde{\mathbf{y}}^{\max} | D_t) - H(\tilde{\mathbf{y}}^{\max} | D_t \cup \{X_t, Y_t\}) \quad (8)$$

$$= \frac{1}{2} \log(2\pi\sigma_t^2(\mathbf{x}_t^{\text{ucb}})) - \frac{1}{2} \log(2\pi\sigma_{t+1}^2(\mathbf{x}_t^{\text{ucb}})) \quad (9)$$

We propose to use Eq. (9) as the surrogate of the mutual information in (3). Since σ_t^2 is independent of $\{X_t, Y_t\}$, the first term in (9) is independent of X_t and can be omitted, thus to maximize (9) in X_t is to minimize $\frac{1}{2} \log(2\pi\sigma_{t+1}^2(\mathbf{x}_t^{\text{ucb}}))$, or equivalently, $\sigma_{t+1}^2(\mathbf{x}_t^{\text{ucb}})$. The explicit expression of $\sigma_{t+1}^2(\mathbf{x}_t^{\text{ucb}})$ is given by the following proposition.

Proposition 1 (Predicted change of GP posterior variance). The posterior variances at any point $\mathbf{x}$ at times t and $t+1$ are related by

$$\sigma_{t+1}^2(\mathbf{x}) = \sigma_t^2(\mathbf{x}) - \gamma(X_t, \mathbf{x}) \quad (10)$$

where

$$\gamma(X_t, \mathbf{x}) \equiv \Sigma_t(\mathbf{x}, X_t) (\Sigma_t(X_t, X_t) + \sigma_0^2 I_m)^{-1} \Sigma_t(X_t, \mathbf{x}) \quad (11)$$

Here, I_m is the m -dimensional identity matrix, $\Sigma_t(\mathbf{x}, X_t)$ stands for the row vector

$$\Sigma_t(\mathbf{x}, X_t) \equiv [\Sigma_t(\mathbf{x}, \mathbf{x}_t^1), \Sigma_t(\mathbf{x}, \mathbf{x}_t^2), \dots, \Sigma_t(\mathbf{x}, \mathbf{x}_t^m)],$$

while $\Sigma_t(\mathbf{x}, X_t) = \Sigma_t(X_t, \mathbf{x})^\top$, and $\Sigma_t(X_t, X_t) = [\Sigma_t(\mathbf{x}, \mathbf{x}')_{\mathbf{x}, \mathbf{x}' \in X_t}]$ is a $m \times m$ matrix.

Proof: Denote I_k as the k -dimensional identity matrix for

any positive integer k , then we have

$$\begin{aligned}\hat{\sigma}_{t+1}^2(\mathbf{x}) &= k(\mathbf{x}, \mathbf{x}) - \\ &L_t(\mathbf{x}) \begin{bmatrix} \mathbf{K}_t + \sigma_0^2 I_m & \mathbf{k}_t(X_t) \\ \mathbf{k}_t(X_t)^\top & k(X_t, X_t) + \sigma_0^2 I_m \end{bmatrix}^{-1} L_t(\mathbf{x})^\top \\ &\text{(using matrix inversion lemma)} \\ &= k(\mathbf{x}, \mathbf{x}) - \mathbf{k}_t(\mathbf{x})^\top (\mathbf{K}_t + \sigma_0^2 I_m)^{-1} \mathbf{k}_t(\mathbf{x}) \\ &\quad - M_t(\mathbf{x}) S_t^{-1} M_t(\mathbf{x})^\top \\ &= \sigma_t^2(\mathbf{x}) - \Sigma_t(\mathbf{x}, X_t) (\Sigma_t^{-1}(X_t, X_t) + \sigma_0^2 I_m) \Sigma_t(X_t, \mathbf{x})\end{aligned}$$

where $L_t(\mathbf{x}) = [\mathbf{k}_t(\mathbf{x}) k(\mathbf{x}, X_t)]$, $M_t(\mathbf{x}) = k(\mathbf{x}, X_t) - \mathbf{k}_t(\mathbf{x})^\top (\mathbf{K}_t + \sigma_0^2 I_m)^{-1} \mathbf{k}_t(X_t)$, $S_t = k(X_t, X_t) + \sigma_0^2 I_m - \mathbf{k}_t(X_t)^\top (\mathbf{K}_t + \sigma_0^2 I_m)^{-1} \mathbf{k}_t(X_t)$. ■

Proposition 1 implies that to maximize the surrogate mutual information (9) in X_t is to maximize $\gamma(X_t, \mathbf{x}_t^{\text{ucb}})$. Note that Eq. (11) does not involve the observation Y_t , which is promising since it avoids the costly sampling over the m -dimensional distribution $Y_t | D_t$ to estimate Y_t as in previous entropy search methods [12], [24]. Unlike other sequential query generation algorithms (for example, GP batch upper confidential bound [17]), we do not need to calculate GP regression (1) m times in each BO iteration.

The resulting multi-agent algorithm, which uses $\gamma(\cdot, \mathbf{x}_t^{\text{ucb}})$ as the acquisition function to calculate batch queries X_t , is listed in Algorithm 1. There is a central coordinator with which all agents communicate. The central coordinator receives the observations from multiple agents, updates the GP model, calculates the queries by maximizing the mutual information according to Eq. (9), and publishes the queries back to the agents. Note that Algorithm 1 can be reduced to a single-agent algorithm without further changes.

Remark 1. It is worth discussing how our algorithm addresses the exploration-exploitation trade-off. In multi-agent BO, this trade-off can be balanced on two dimensions: time and batch. In our algorithm, the batch-dimension trade-off is balanced through the design of γ while the time-dimension trade-off is guided by $\mathbf{x}_t^{\text{ucb}}$.

The time-dimension trade-off is that given the observations in the past, the next query points should balance between visiting the empirically good locations (exploitation) and covering under-explored locations (exploration). This trade-off is already studied intensively in single-agent BO. One of the most notable methods is to use UCB as the acquisition function, like the one in our definition of $\mathbf{x}_t^{\text{ucb}}$. These algorithms typically favor locations with high values in the sum of the exploitation term μ_t and the exploration term σ_t .

The batch-dimension trade-off arises only in the multi-agent setting. It means the query points at the same iteration should remain close to some empirically high-value locations and be sufficiently diverse. This trade-off can be seen through the tension between the exploitation term $\Sigma_t(\mathbf{x}, X_t)$ (and its transpose) and the exploration term $(\Sigma_t(X_t, X_t) + \sigma_0^2 I_m)^{-1}$ in $\gamma(X_t, \mathbf{x}_t^{\text{ucb}})$. Without loss of generality, let us only consider the magnitudes of these terms. With $\gamma(X_t, \mathbf{x}_t^{\text{ucb}})$ as the acquisition function, X_t should be selected so that $\Sigma_t(X_t, \mathbf{x}_t^{\text{ucb}})$ is large, thus the query points X_t should be highly correlated with $\mathbf{x}_t^{\text{ucb}}$, which typically implies $\mathbf{x}_t^1, \mathbf{x}_t^2, \dots, \mathbf{x}_t^m$ are spatially close to $\mathbf{x}_t^{\text{ucb}}$.

Algorithm 1: Multi-Agent Gaussian Max-value Entropy Search (GMES)

Initialization: Agent number m , Maximal iterations T , Observation noise σ_0^2 , Confidence interval width β_t , A central coordinator with Data set $D_0 = \emptyset$, Initial queries X_0 , Gaussian process prior model $GP(\mu_0, \Sigma_0)$: μ_0 is the zero function, Σ_0 is the white kernel with noise level σ_0^2 .

Output: Inferred maximum, $\arg\max_{\mathbf{x} \in \mathcal{X}} \mu_T(\mathbf{x})$, from $GP(\mu_T, \Sigma_T | D_T)$.

```

1 for  $t = 1, 2, \dots, T$  do
2   # Each agent gets its observation
   at the query point
3   for Agent  $i = 1, 2, \dots, m$  do
4     Observe  $\mathbf{y}_{t-1}^i = f(\mathbf{x}_{t-1}^i) + \epsilon$ ,  $\epsilon \sim \mathcal{N}(0, \sigma_0^2)$ 
5   # All agents return the
   observation to the central
   coordinator
6    $D_t = D_{t-1} \cup \{X_{t-1}, Y_{t-1}\}$ , where
    $Y_{t-1} = \{\mathbf{y}_{t-1}^1, \mathbf{y}_{t-1}^2, \dots, \mathbf{y}_{t-1}^m\}$ 
7   # Central coordinator calculates
   queries
8    $\mathbf{x}_t^{\text{ucb}} \leftarrow \arg\max_{\mathbf{x} \in \mathcal{X}} \mu_t(\mathbf{x}) + \beta_t \sigma_t(\mathbf{x})$ 
9    $X_t = \{\mathbf{x}_t^1, \mathbf{x}_t^2, \dots, \mathbf{x}_t^m\} \leftarrow$ 
    $\arg\max_{X \in \mathcal{X}^m} \gamma(X, \mathbf{x}_t^{\text{ucb}})$ 
10  # Central coordinator publishes
   queries to agents
11  Publish queries  $\mathbf{x}_t^i$  to agent  $i$  for  $i \in \{1, \dots, m\}$ 
```

Meanwhile, X_t should also be chosen such that $\Sigma_t(X_t, X_t)$ is close to the zero matrix, meaning the correlation between the points in X_t themselves are small, which typically implies $\mathbf{x}_t^1, \mathbf{x}_t^2, \dots, \mathbf{x}_t^m$ maintain some spatial separation among themselves. The two components of γ thus balance the tendency of the query points to stay close to $\mathbf{x}_t^{\text{ucb}}$ and to maintain spatial separation among themselves, respectively.

B. Practical Implementation of Algorithm 1

A few details of Algorithm 1 need to be modified for efficient computation and safety considerations in the real world. We briefly discuss the changes below. The full version of the resulting algorithm can be found in Appendix B of our online report [29].

1) *Calculating Queries by Gradient Ascent:* One computational challenge of Algorithm 1 is to find the $X \in \mathcal{X}^m$ that maximizes $\gamma(X, \mathbf{x}_t^{\text{ucb}})$. Brute-force methods such as grid or random search are unsuitable for multi-agent implementations as the search space grows exponentially in the number of agents m . We mitigate the expensive computation in brute-force methods and approximate the maximization of γ through gradient ascent. Given the analytical form of γ , its gradients can be computed efficiently using standard auto-differentiation software packages like PyTorch [30].

We randomly initialize the batched queries for gradient ascent. To ensure the convergence of the gradient updates, an appropriate sequence of step sizes $\{\delta_t \mid t = 1, 2, \dots, T\}$ needs to be applied to the gradient. In practice, algorithms such as Adam [31] would suffice to decide the step sizes.

Finally, to guarantee that the algorithm returns a X_t that is inside the compact domain $\mathcal{X}^m$, we project the query points back to $\mathcal{X}^m$ after every gradient update. The projection operator is defined by

$$\Gamma_{\mathcal{X}^m}(X) = \operatorname{argmin}_{\bar{X} \in \mathcal{X}^m} \|\bar{X} - X\|_2. \quad (12)$$

And line 10 in Algorithm 1 can be replaced with multiple iterations of the gradient ascent update described below.

$$X_t \leftarrow \Gamma_{\mathcal{X}^m}(X_t + \delta_t \cdot \nabla_{X_t} \gamma(X_t, \mathbf{x}_t^{\text{ucb}}))$$

2) *Ensuring Safety with Log-Barrier*: Safety considerations, like collision avoidance, are common in many real-world multi-agent exploration tasks. Safety considerations require query points in the same batch X_t to be separated by at least the physical size of the robot. When this constraint is enforced, we subtract a log-barrier term $p(X_t)$ from the acquisition function in (11). The log-barrier $p(X_t)$ is defined by

$$p(X_t) = \sum_{i,j:1 \leq i < j \leq m} \left[-\frac{1}{L} \log(\|\mathbf{x}_t^i - \mathbf{x}_t^j\| - r_{\text{div}}) \right]^+ \quad (13)$$

where $[\cdot]^+$ means projection to the positive half-space $[0, \infty)$. The parameter r_{div} specifies the minimal separation between the robots. We initialize X_t to ensure $\|\mathbf{x}_t^i - \mathbf{x}_t^j\| > r_{\text{div}}$ for all $i \neq j$, so that $p(X_t)$ is well-defined in all iterations of the gradient updates.

IV. NUMERICAL EXPERIMENTS

We conduct numerical experiments with standard test functions to show the advantage of our proposed algorithm compared to recent multi-agent BO baselines. The open-source implementation of the numerical experiments can be found on <https://github.com/mahaitongdae/dbo>.

A. Numerical Experiment Setup

We first conduct numerical experiments on several standard test functions. Figure 2 shows the optimization landscapes of the test functions for the numerical experiments. The Ackley function has many local maxima, but only one global maximum, and its value tends to be higher at points closer to the origin. The Bird function has two global maxima and two local maxima. The Rosenbrock function has a global maximum surrounded by many saddle points. We use the Matérn kernel as the kernel function $k(\cdot, \cdot)$ in (1) [26]. The log-barrier term is not applied to our algorithm in the numerical experiments. The algorithm parameters are listed in Table I.

The algorithms are evaluated under two performance metrics, instant regret R_t and cumulative regret $\bar{R}_t$, defined by

$$R_t = f^* - \max_{\tau \in \{1, 2, \dots, t\}} \max_{\mathbf{x}_\tau \in X_\tau} f(\mathbf{x}_\tau), \quad \bar{R}_t = \sum_{\tau=0}^t R_\tau \quad (14)$$

TABLE I: Algorithm parameters for the numerical experiments

Notation	Meaning	Value
T	Total iterations	150
σ_0	Observation noise	0.1
N	Gradient ascent iterations	50
β_t	Confidence interval width	$3 - 0.01 \cdot t$
ν	Matérn kernel parameter	1.5
δ_t	Step sizes for gradient update	Adam [31]
$p(\cdot)$	Log-barrier term	Not applied

where $f^* = \max_{\mathbf{x} \in \mathcal{X}} f(\mathbf{x})$. Note that the instant regret R_t only considers the best observation among all agents over all iterations, which is a fair metric even for those algorithms that explore aggressively (like the UCB-PE).

B. Experiment Results for multi-agent BO

We compare our algorithm with several baseline multi-agent BO algorithms from recent literature, including GP-UCB with pure exploitation (GP-UCB-PE) [15], Gaussian process batch upper confidence bound (GP-BUCB) [17], EI with Monte-Carlo sampling (EI-MCMC) [18], EI with stochastic policies (EI-SP) [14], and Parallel Thompson sampling (TS) [13], [15].

The initial query points are randomly selected using a fixed random seed in each experiment. We run each algorithm five times and report the instant and cumulative regrets for 5-agent experiments in Figure 1. Our proposed algorithm, labeled ES, has the lowest instant regret on all test functions with five agents. It also has the lowest cumulative regret on the Ackley and Bird functions. EI-SP and UCB-PE achieve better cumulative regret than our algorithm on the Rosenbrock function, which shows that our algorithm might take more iterations to reach its lowest instant regret than these two baselines in this task. Table II shows the instant regret at the final iteration for 10- and 30-agent experiments. Our algorithm is still consistently better than the baseline algorithms in instant regrets, being the second-best on the Ackley and Rosenbrock tasks in the 30-agent experiment and the best for all other tasks. Even for those tasks where ES is the second-best, its regret is rather close to being the lowest one. More experimental results, including further comparing different numbers of agents up to 50, could be found in Appendix D in the online report [29].

C. Analysis of Query Distributions

The following experiments demonstrate our algorithm's advantage in adapting the exploration-exploitation trade-off to different test functions. Figure 3 plots the query point distributions of the proposed algorithm with gradient ascent and two baselines, BUCB and UCB-PE, that assign fixed exploration/exploitation roles to the agents. The data comes from the 10-agent experiments on Ackley and Rosenbrock tasks in the previous subsection. Our proposed algorithm, ES, has shown the ability to adjust its exploration-exploitation balance for different tasks. It exploits the global landscape of Ackley by clustering its query points near the origin, where the function values are generally higher. It also generates a diverse query distribution on Rosenbrock, allowing its inferred maximum (i.e., $\operatorname{argmax}_{\mathbf{x} \in \mathcal{X}} \mu_t(\mathbf{x})$) to approach the




Fig. 1: Instant regret (upper row) and cumulative regret (lower row) with five agents. The solid lines are the average performance, and the shaded regions are the 95 percent confidence interval across five runs. Our algorithm is colored red with label ES.

TABLE II: Mean and Variance of instant regret($\times 10^{-2}$) at the last iteration (150 iterations in total) for 10- and 30-agent experiments.

	ES-10 (ours)	EI-SP-10	GP-BUCB-10	GP-UCBPE-10	TS-10	EI-MCMC-10
Ackley	3.383	52.26	3.411	4.619	5.636	19.36
	± 0.503	± 12.26	± 0.800	± 1.00	± 2.243	± 4.259
Bird	3.626	33.205	5.057	2.671	5.632	3.954
	± 1.409	± 11.696	± 2.178	± 0.639	± 1.955	± 1.038
Rosenbrock	1.030	11.08	21.12	1.572	41.56	1.283
	± 0.313	± 3.153	± 10.52	± 0.264	± 7.724	± 0.551
	ES-30 (ours)	EI-SP-30	GP-BUCB-30	GP-UCBPE-30	TS-30	EI-MCMC-30
Ackley	3.279	3.570	4.670	3.411	3.218	12.047
	± 2.604	± 2.300	± 2.805	± 2.282	± 2.089	± 6.094
Bird	1.857	2.026	2.244	2.026	2.421	2.149
	± 1.433	± 1.569	± 1.990	± 1.569	± 2.365	± 1.803
Rosenbrock	0.859	2.370	49.653	1.540	0.436	1.447
	± 1.121	± 1.990	± 24.728	± 0.659	± 0.847	± 2.994




Fig. 2: Test Functions in the Numerical Experiments. From left to right: Ackley, Bird, Rosenbrock.

true maximum as more observations are made. The final instant regrets of ES are low on both tasks.

In contrast, the baseline algorithms do not adapt sufficiently to the two tasks. UCB-PE always has sparse query distributions, which allows it to explore the relatively flat landscape of Rosenbrock and gives it decent performance on this task; however, UCB-PE is not the best algorithm on Ackley since it does not sufficiently exploit the prominent peak at the origin. BUCB outperforms UCB-PE on Ackley since it exploits well. However, BUCB has little diversity in its queries on Rosenbrock, making its exploration insufficient and its instant regret much higher than the other two algorithms. The inferred maximum of BUCB also gets stuck in a saddle point instead of converging to the true maximum. This lack of diversity of BUCB's query points on Rosenbrock could be due to BUCB's iterative greedy approach to deciding the batch of query points.

The inflexibility of BUCB and UCB-PE is ultimately

caused by their fixed exploration/exploitation role assignment regardless of the test function. In comparison, our algorithm selects query points that return the most informative queries for different tasks by maximizing the acquisition function $\gamma(\cdot, \mathbf{x}_t^{\text{ucb}})$. Our algorithm can thus adapt its exploration-exploitation balance to the test functions above and outperform BUCB and UCB-PE.

V. SOURCE SEEKING PHYSICAL EXPERIMENTS

We demonstrate the effectiveness of our algorithm on real robots through experiments on the multi-agent source seeking problem. The task is to let multiple robotic vehicles with sensing abilities collaborate to locate a source of interest. Examples of source seeking problems include pollutant source localization [32], distributed sensor placement [33], target tracking [34], and so on. Typically, the source location is where the sensor reading is the strongest; therefore, the problem is often viewed as equivalent to driving the robots to maximize their sensor readings. Measuring the source signal is usually inexpensive, but driving the robots to desired locations could be energy- and time-consuming. Therefore, locating the source with limited samples is desirable, and BO algorithms are suitable for source seeking applications due to their sample efficiency. The full version of our experiment video can be found on <https://youtu.be/>




Fig. 3: Query Distributions. The four columns on the left are the results of the Ackley task; the other four are the results of the Rosenbrock task. Every row indicates all the queries up to the 10, 50, 100, and 150 iterations. In each figure: the background contour plot indicates the true objective function value; the orange star marks the global maximum; the yellow square marks the inferred maximum (the maximum location of the posterior mean function μ_t) at the corresponding iteration; green dots are queries generated the algorithms; the title indicates the instant regret at the final iteration.

PK_emQ85sb0.

A. Experiment Setup

In the experiments, we use four ROBOTIS TurtleBot3 with onboard light sensors, as illustrated in Figure 4(d). The goal is to find the location of the highest brightness on the ground level in a dark room. We use a desktop computer as the central coordinator to maintain the GP model and calculate the queries. Figure 4 shows three different source seeking experimental setups. The simplest task, SINGLE, has only one LED lamp hung above the ground, corresponding to the brightest location in the room. The other two tasks, labeled SPARSE and DENSE, are with four lamps in the room, where two of the lamps are brighter and the other two dimmer. Each lamp is hung at a different height, and the brighter lamp hanging closest to the ground, marked by red boxes in Figures 4(b) and 4(c), corresponds to the brightest location in the room. The only differences between the two four-light experiments are the lamp spacing.

We use the basic look-ahead PID controller to drive the Turtlebots to the query points our multi-agent BO algorithm decides while avoiding collision between the robots. See Appendix C of our online report [29] for details about our robot controller. We include the log-barrier safety term in the acquisition function to ensure the robots' target positions do not induce collisions. The experiments terminate when the inferred maximum (the maximum location of the posterior mean function μ_t) is within 0.1m of the actual brightest location for three consecutive iterations.

B. Results

The accompanying video shows that the robots consistently find the highest brightness location within a short




Fig. 4: Source-seeking experiment setups for different tasks. In each of Figures 4(a)-4(c): the largest subfigure provides the side and top views of the experiment, where the target lamp to seek is marked by a red box; the upper right illustrates the relative location of the robots (blue circles) and the lights (red, sun-shaped icons) projected to the ground plane; the lower right shows the contour plot about the current GP posterior mean, where red and blue colors indicate high and low posterior mean values, respectively. Figure 4(d) illustrates the light sensor setup on the robot.

time. The results suggest our algorithm can be applied to a multi-agent team and efficiently maximize a general black-box function in the real world. Figure 5 further shows the performance difference of source seeking with four robots compared to using only one robot. The multi-agent team saves the source seeking time by 59.9% and the iterations by 67.6% compared with the single agent. This impressive advantage in efficiency over the single-agent approach shows the multi-agent BO approach can significantly benefit time-critical source seeking applications, such as search-and-




Fig. 5: Comparisons between the single- and multi-agent approaches regarding the total seeking time (in seconds), the sum of distance traveled, and total iterations. The x-axis labels indicate the environment setups, as described in Figure 4.

rescue missions. More experimental results, like the robot trajectories, could be found in Appendix D in the online report [29].

VI. CONCLUSION

In this paper, we proposed the Gaussian max-value Entropy Search (GMES), a computationally efficient algorithm for multi-agent Bayesian optimization. We use the normal distribution to approximate the posterior function max-value to design an acquisition function that balances the exploration-exploitation trade-off and has a closed-form expression that allows simple computation without complex approximations and relaxations. We further improve the algorithm by using gradient ascent for the query point optimization and a log-barrier term to enforce safety constraints. Experiment results show that the GMES outperforms other multi-agent BO baselines in the numerical experiments and effectively seeks light sources on real robots. Future works include analyzing the performance of GMES theoretically and using GMES with the distributed multi-agent settings with local communications only.

REFERENCES

- [1] F. Berkenkamp, A. P. Schoellig, and A. Krause, "Safe controller optimization for quadrotors with gaussian processes," in *2016 IEEE International Conference on Robotics and Automation (ICRA)*, 2016, pp. 491–496.
- [2] W. Zhao, T. He, and C. Liu, "Model-free safe control for zero-violation reinforcement learning," in *Proceedings of the 5th Conference on Robot Learning*, ser. Proceedings of Machine Learning Research, A. Faust, D. Hsu, and G. Neumann, Eds., vol. 164. PMLR, 08–11 Nov 2022, pp. 784–793.
- [3] H. Ma, C. Liu, S. E. Li, S. Zheng, and J. Chen, "Joint synthesis of safety certificate and safe control policy using constrained reinforcement learning," in *Proceedings of The 4th Annual Learning for Dynamics and Control Conference*, ser. Proceedings of Machine Learning Research, vol. 168. PMLR, 23–24 Jun 2022, pp. 97–109.
- [4] T. Zhang, V. Qin, Y. Tang, and N. Li, "Source seeking by dynamic source location estimation," in *2021 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2021, pp. 2598–2605.
- [5] —, "Distributed information-based source seeking," *arXiv preprint arXiv:2209.09421*, 2022.
- [6] K. Rosser, J. Kok, J. Chahl, and J. Bongard, "Sim2real gap is non-monotonic with robot complexity for morphology-in-the-loop flapping wing design," in *2020 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2020, pp. 7001–7007.
- [7] P. Auer, "Using confidence bounds for exploitation-exploration trade-offs," *Journal of Machine Learning Research*, vol. 3, no. Nov, pp. 397–422, 2002.
- [8] N. Srinivas, A. Krause, S. M. Kakade, and M. Seeger, "Gaussian process optimization in the bandit setting: No regret and experimental design," *arXiv preprint arXiv:0912.3995*, 2009.
- [9] H. J. Kushner, "A New Method of Locating the Maximum Point of an Arbitrary Multipipe Curve in the Presence of Noise," *Journal of Basic Engineering*, vol. 86, no. 1, pp. 97–106, Mar. 1964.
- [10] J. Moćkus, "On bayesian methods for seeking the extremum," in *Optimization techniques IFIP technical conference*. Springer, 1975, pp. 400–404.
- [11] J. Wu, M. Poloczek, A. G. Wilson, and P. Frazier, "Bayesian optimization with gradients," *Advances in neural information processing systems*, vol. 30, 2017.
- [12] P. Hennig and C. J. Schuler, "Entropy Search for Information-Efficient Global Optimization," *Journal of Machine Learning Research*, p. 29, Jun. 2012.
- [13] W. R. Thompson, "On the likelihood that one unknown probability exceeds another in view of the evidence of two samples," *Biometrika*, vol. 25, no. 3–4, pp. 285–294, 1933.
- [14] J. Garcia-Barcos and R. Martinez-Cantin, "Fully distributed bayesian optimization with stochastic policies," *arXiv preprint arXiv:1902.09992*, 2019.
- [15] E. Contal, D. Buffoni, A. Robicquet, and N. Vayatis, "Parallel gaussian process optimization with upper confidence bound and pure exploration," in *Joint European Conference on Machine Learning and Knowledge Discovery in Databases*. Springer, 2013, pp. 225–240.
- [16] B. Shahriari, A. Bouchard-Côté, and N. Freitas, "Unbounded bayesian optimization via regularization," in *Artificial intelligence and statistics*. PMLR, 2016, pp. 1168–1176.
- [17] T. Desautels, A. Krause, and J. W. Burdick, "Parallelizing exploration-exploitation tradeoffs in gaussian process bandit optimization," *Journal of Machine Learning Research*, vol. 15, pp. 3873–3923, 2014.
- [18] J. Snoek, H. Larochelle, and R. P. Adams, "Practical bayesian optimization of machine learning algorithms," *Advances in neural information processing systems*, vol. 25, 2012.
- [19] K. Kandasamy, A. Krishnamurthy, J. Schneider, and B. Póczos, "Parallelised bayesian optimisation via thompson sampling," in *International Conference on Artificial Intelligence and Statistics*. PMLR, 2018, pp. 133–142.
- [20] A. Shah and Z. Ghahramani, "Parallel predictive entropy search for batch global optimization of expensive objective functions," *Advances in neural information processing systems*, vol. 28, 2015.
- [21] J. Wu and P. Frazier, "The parallel knowledge gradient method for batch bayesian optimization," *Advances in neural information processing systems*, vol. 29, 2016.
- [22] P. Hennig, "Optimal reinforcement learning for gaussian systems," *Advances in Neural Information Processing Systems*, vol. 24, 2011.
- [23] Z. Wang and S. Jegelka, "Max-value Entropy Search for Efficient Bayesian Optimization," in *Proceedings of the 34th International Conference on Machine Learning*. PMLR, Jul. 2017, pp. 3627–3635, iSSN: 2640-3498.
- [24] J. M. Hernández-Lobato, M. W. Hoffman, and Z. Ghahramani, "Predictive Entropy Search for Efficient Global Optimization of Black-box Functions," in *Advances in Neural Information Processing Systems*, vol. 27. Curran Associates, Inc., 2014.
- [25] B. Tu, A. Gandy, N. Kantas, and B. Shafei, "Joint Entropy Search for Multi-objective Bayesian Optimization," Oct. 2022, arXiv:2210.02905 [cs, math, stat].
- [26] C. K. Williams and C. E. Rasmussen, *Gaussian processes for machine learning*. MIT press Cambridge, MA, 2006, vol. 2, no. 3.
- [27] T. M. Cover, *Elements of information theory*. John Wiley & Sons, 1999.
- [28] M. Seeger, "Gaussian processes for machine learning," *International journal of neural systems*, vol. 14, no. 02, pp. 69–106, 2004.
- [29] H. Ma, T. Zhang, Y. Wu, F. Calmon, and N. Li, "Gaussian max-value entropy search for multi-agent bayesian optimization," 2023. [Online]. Available: https://scholar.harvard.edu/files/haitongma/files/gaussian_max_value_entropy_search_online_report.pdf
- [30] J. Gardner, G. Pleiss, K. Q. Weinberger, D. Bindel, and A. G. Wilson, "Gpytorch: Blackbox matrix-matrix gaussian process inference with gpu acceleration," *Advances in neural information processing systems*, vol. 31, 2018.
- [31] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," *arXiv preprint arXiv:1412.6980*, 2014.
- [32] B. Bayat, N. Crasta, H. Li, and A. Ijspeert, "Optimal search strategies for pollutant source localization," in *2016 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. Ieee, 2016, pp. 1801–1807.
- [33] R. Bachmayer and N. E. Leonard, "Vehicle networks for gradient descent in a sampled environment," in *Proceedings of the 41st IEEE Conference on Decision and Control*, 2002., vol. 1. IEEE, 2002, pp. 112–117.
- [34] F. Morbidi and G. L. Mariottini, "Active target tracking and cooperative localization for teams of aerial vehicles," *IEEE transactions on control systems technology*, vol. 21, no. 5, pp. 1694–1707, 2012.