

2020-07

Optimal composition of heterogeneous multi-agent teams for coverage problems with performance bound guarantees

Chuangchuang Sun, Shirantha Welikala, Christos G Cassandras. 2020. "Optimal composition of heterogeneous multi-agent teams for coverage problems with performance bound guarantees." *Automatica*, Volume 117, pp. 108961 - 108961. <https://doi.org/10.1016/j.automatica.2020.108961>
<https://hdl.handle.net/2144/42871>

"Downloaded from OpenBU. Boston University's institutional repository."

Optimal Composition of Heterogeneous Multi-Agent Teams for Coverage Problems with Performance Bound Guarantees^{*}

Chuangchuang Sun^a, Shirantha Welikala and Christos G. Cassandras^b

^aDepartment of Aeronautics and Astronautics, Massachusetts Institute of Technology, Cambridge, MA 02139, USA.

^bDivision of Systems Engineering and Center for Information and Systems Engineering, Boston University, Brookline, MA 02446, USA.

Abstract

We consider the problem of determining the optimal composition of a heterogeneous multi-agent team for coverage problems by including costs associated with different agents and subject to an upper bound on the maximal allowable number of agents. We formulate a resource allocation problem without introducing additional non-convexities to the original problem. We develop a distributed Projected Gradient Ascent (PGA) algorithm to solve the optimal team composition problem. To deal with non-convexity, we initialize the algorithm using a greedy method and exploit the submodularity and curvature properties of the coverage objective function to derive novel tighter performance bound guarantees on the optimization problem solution. Numerical examples are included to validate the effectiveness of this approach in diverse mission space configurations and different heterogeneous multi-agent collections. Comparative results obtained using a commercial mixed-integer nonlinear programming problem solver demonstrate both the accuracy and computational efficiency of the distributed PGA algorithm.

Key words: Multi-agent Systems, Optimization, Cooperative Control,

1 Introduction

Cooperative multi-agent systems are pervasive in a number of applications, including but not limited to, surveillance [7, 22], search and rescue missions [15], consensus [26, 27] and agriculture [1]. One of the most basic tasks such a system can perform that has seen a wide range of applications is *coverage*. The fundamental multi-agent optimal coverage problem has been extensively studied in the literature, e.g., [3–6, 16]. In this problem, agents are deployed to “cover” as much of a given mission space as possible in the sense that the team aims to optimally jointly detect events of interest (e.g., data sources) that may randomly occur anywhere in this space. The coverage performance is measured by an appropriate metric, which is normally defined as the joint event detection probability. The optimal coverage problem is particularly challenging due to the generally non-

convex nature of this metric and the non-convexity of the mission space itself due to the presence of obstacles which act as constraints on the feasible agent locations that constitute a solution to the problem.

Thus far, the analysis of the optimal coverage problem has been carried out based on the assumption that there exists a fixed number N of agents to be deployed. However, this number is often limited by cost constraints, leading to a natural trade-off between coverage performance (which is normally monotonically increasing in N) and total system cost. In such a setting, an additional aspect of the problem is that of managing a set of *heterogeneous* agents: when agents fall into different classes characterized by different properties such as sensing capacity, range, attenuation rate, and cost, then the problem becomes one of determining the *optimal cooperative team composition* in terms of the number of agents selected from each class so as to optimize an appropriate metric capturing the performance-cost trade-off. Clearly, it is possible that a certain team composition can achieve the same coverage performance as another, but with a lower cost due to the heterogeneity of agents. The purpose of this paper is to address the optimal coverage problem in the presence of heterogeneous agents under cost constraints.

As mentioned above, the optimal coverage problem is already challenging due to its non-convex nature. Heuristic

^{*} This paper was not presented at any conferences. For C. Sun, this work was performed while at Boston University. This work is supported in part by NSF under grants ECCS-1509084, DMS-1664644, CNS-1645681, by AFOSR under grant FA9550-19-1-0158, by ARPA-E’s NEXTCAR program under grant dea0000796 and by the MathWorks.

Email addresses: ccsun1@mit.edu (Chuangchuang Sun), {shiran27, cgc}@bu.edu. (Shirantha Welikala and Christos G. Cassandras).

algorithms (e.g., genetic algorithms [10]), are often used and may lead to empirically near-global optimality, but they are prohibitively inefficient for on-line use. On the other hand, on-line algorithms sacrifice potential optimality to achieve efficiency; this includes distributed gradient-based algorithms [6, 13, 28] and Voronoi-partition-based algorithms [3, 9, 12] which lead to generally locally optimal solutions. Methods for efficiently escaping such local optima using a “boosting function” approach were proposed in [19, 25], while a decentralized control law in [18] seeks a combination of optimal coverage and exploration of the area of interest.

A parallel effort to deal with the difficulty of finding a globally optimal solution for the basic coverage problem is by exploiting the *submodularity* properties of the coverage performance functions used (e.g., the joint event detection probability). This is accomplished in [20] by using a greedy algorithm to initialize the state of the system (i.e., the locations of the agents), followed by a conventional gradient ascent technique to obtain an improved (still locally optimal) solution. Due to submodularity, the ratio f^G/f^* , where f^G and f^* correspond to the objective function values under a greedy solution and the globally optimal solution respectively, has a lower bound $L \leq f^G/f^*$ which is shown to be $L = 1/2$ in [11]. When the objective function f is *monotone submodular* (which applies to coverage metrics), then it has been shown that $L = (1 - \frac{1}{e})$ [17] and becomes $L = (1 - (1 - \frac{1}{N})^N)$ when the allowable maximum number of agents is constrained to N . Recent work [8], [14, 23] has further improved these performance bounds by exploiting the specific nature of the monotonicity (also known as *curvature* properties) of the specific objective function. By using these improved bounds, the solutions to a variety of optimal coverage problems in [20] have been shown to often approach $L = 1$, i.e., to yield almost globally optimally solutions.

Our contributions in this paper are threefold. First, we formulate the problem of determining an optimal team composition under a heterogeneous set of agents as a resource allocation problem without introducing additional non-convexity features to it. In particular, instead of treating the (discrete) number of agents in each class as a decision variable, we associate this number with the (continuous) sensing capacity of the agents in each class; hence, an allocation of zero sensing capacity implies a virtual (or non-existing) agent. In our problem formulation, instead of imposing a hard cardinality constraint, an l_1 norm penalty in the objective function is employed to induce sparsity and prevent any new non-convexity from being introduced.

Secondly, for the coverage component of the objective function (i.e., without the aforementioned penalty term), a greedy algorithm is used and two new improved performance bounds are derived based on the concepts of partial curvature [14], total curvature, and greedy curvature [8].

Finally, we propose a distributed projected gradient ascent

algorithm to solve the overall optimal team composition problem. The key to this algorithm is the proper selection of an *initial condition* which is characterized by a provable lower bound. Thus, we first use a greedy method to generate a candidate solution to the underlying coverage component of the problem which always contains all the available agents. This is used as the initial condition to solve the main problem (combining coverage and system cost). In doing so, a distributed projected gradient ascent scheme is used whose final solution recovers both the integer and real variables associated with the problem which respectively define the optimal team composition and the optimal agent locations.

Relative to our previous work [20, 28], here we consider a significantly different problem and make a number of key contributions to the coverage control problem with heterogeneous agents. A crucial difference in this work compared to both [20, 28] is that we do not assume that a given number of agents is to be deployed; rather, we seek to determine the number of agents (subject to an upper bound constraint) and optimal team composition (not only the optimal agent locations), which is a combinatorial NP-hard problem. Moreover, heterogeneity considered in this work brings challenges to the aforementioned greedy algorithm, to the associated performance bounds, and to the process of determining an optimal team composition - all of which are addressed here. Finally, two new tighter performance bounds are derived compared to those in [20].

The rest of the paper is organized as follows. The optimization problem for determining the optimal team composition is formulated in Section 2. Then, to obtain a good initial condition to solve this optimization problem, a greedy algorithm is presented in Section 3, along with some performance bound guarantees. Subsequently, to completely solve the formulated optimization problem, a distributed projected gradient ascent process is proposed in Section 4, along with some theoretical results regarding the nature of its terminal solution. Numerical results are included in Section 5 to validate the effectiveness of the proposed solution technique. Finally, Section 6 concludes the paper.

1.1 Preliminaries

Some notations used throughout this paper are introduced here. The n -dimensional Euclidean space is denoted by \mathbb{R}^n . Lowercase letters are used to denote vectors (E.g. $x \in \mathbb{R}^n$) and bold (and lowercase) letters are used to denote matrices (E.g. $\mathbf{s} \in \mathbb{R}^{N \times 2}$) while uppercase letters are used to denote set variables. Moreover, $|\cdot|$ and $\|\cdot\|$ denote the cardinality of a set variable and the l_2 norm of a vector respectively.

2 Problem Formulation

We begin with a brief review of the the multi-agent coverage problem (see [5, 9, 28]). The mission space $\Omega \subseteq \mathbb{R}^2$ is modeled as a convex compact polygon. For non-convex

polygons Ω_1 , such as the self-intersecting ones, we make Ω the convex hull of Ω_1 , while $\Omega \setminus \Omega_1$ defines obstacles that agents have to avoid. Let $R(x) : \mathbb{R}^2 \rightarrow \mathbb{R}$ be an event density function such that $R(x) \geq 0, \forall x \in \Omega$ and $\int_{\Omega} R(x) dx < \infty$ such that $R(x)$ represents the relative importance of a point $x \in \Omega$. Obstacles present in the mission space can both limit the movement of agents and interfere with their sensing capacities. Such obstacles are modeled as non-intersecting polygons M_1, \dots, M_m and their interiors are forbidden regions for the agents. As a result, the feasible (safety) region is $F = \Omega \setminus (\overset{\circ}{M}_1 \cup \dots \cup \overset{\circ}{M}_m)$, where $\overset{\circ}{M}$ is the interior of M .

With N as the maximum possible number of agents, we have $\mathbf{s} = [s_1^T, \dots, s_N^T]^T \in \mathbb{R}^{N \times 2}$ denoting the locations of the N agents with each $s_i \in \mathbb{R}^2, \forall i = 1, \dots, N$. Then, the following sensing model is adopted. For any point $x \in \Omega$ and a certain agent at s_i , there are two issues affecting if the agent can detect an event occurring at x . First, the agent is characterized by a sensing region defined as $\Omega_i = \{x \mid \|x - s_i\| \leq \delta_i\}$, where δ_i is the sensing range. Secondly, obstacles prevent a signal at x from reaching s_i . This is described by the condition $\eta s_i + (1 - \eta)x \in F, \eta \in [0, 1]$, i.e., the segment connecting x and s_i must be contained in the feasible region. Then, the visibility set of s_i is defined as $V(s_i) = \Omega_i \cap \{x \mid \eta s_i + (1 - \eta)x \in F\}$ and the invisibility set $\bar{V}(s_i)$ is the complement of $V(s_i)$ in F , i.e., $\bar{V}(s_i) = F \setminus V(s_i)$. An illustration of $V(s_i)$ is shown in Fig. 1.

The probability that agent i detects an event at x in an unconstrained environment is given by

$$p_i(x, s_i) = p_{i0} e^{-\lambda_i \|x - s_i\|} \quad (1)$$

where $p_{i0} \in (0, 1]$ is the agent's sensing capacity and $\lambda_i > 0$ is a sensing decay (attenuation) factor. As discussed in the introduction, different $p_i(x, s_i)$ specified by p_{i0}, δ_i and λ_i will lead to a heterogeneous multi-agent system. In a mis-

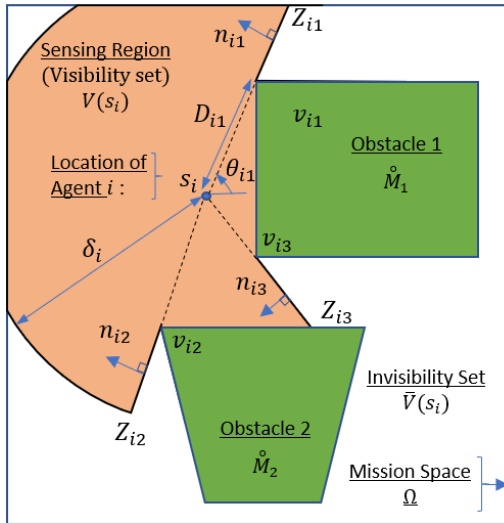


Fig. 1. Mission space with obstacles.

sion space with constraints, the agent's detection probability becomes:

$$\hat{p}_i(x, s_i) = \begin{cases} p_i(x, s_i) & \text{if } x \in V(s_i), \\ 0 & \text{otherwise.} \end{cases} \quad (2)$$

Finally, assuming detection independence among the N agents, the joint detection probability of an event at x is given by

$$\hat{P}(x, \mathbf{s}) = 1 - \prod_{i=1}^N (1 - \hat{p}_i(x, s_i)).$$

As formulated in [28], the optimal multi-agent coverage problem is

$$\begin{aligned} \max_{\mathbf{s}} H(\mathbf{s}) &= \int_{\Omega} R(x) \hat{P}(x, \mathbf{s}) dx \\ \text{s.t. } & s_i \in F, \quad i = 1, \dots, N, \end{aligned} \quad (3)$$

where the number of the agents N is a predetermined constant. When N is in fact an additional decision variable constrained by the cost of agents, we proceed by capturing the trade-off between improved performance, which monotonically increases with N , and agent cost as follows. Letting N be the the *maximum possible number* of agents to consider, we formulate a resource (sensing capacity) allocation problem:

$$\begin{aligned} \max_{\mathbf{s}, \mathbf{t}} H(\mathbf{s}, \mathbf{t}) &= \int_{\Omega} R(x) P(x, \mathbf{s}, \mathbf{t}) dx - \beta \sum_{i=1}^N t_i \\ \text{s.t. } & s_i \in F, \quad t_i \in \{0, 1\}, \quad i = 1, \dots, N, \end{aligned} \quad (4)$$

with,

$$P(x, \mathbf{s}, \mathbf{t}) = 1 - \prod_{i=1}^N (1 - t_i \hat{p}_i(x, s_i)). \quad (5)$$

In (4), $\mathbf{t} = [t_1, t_2, \dots, t_N]^T$ and t_i is a binary decision variable associated with agent i . The term $\beta \sum_{i=1}^N t_i$ denotes the cost of deploying N agents, where $\beta \geq 0$ is a weight capturing the cost of each agent (assumed to be the same in this formulation). In order to ensure a properly normalized objective function, β must be selected to be consistent with the following convex combination of objectives:

$$\tilde{H}(\mathbf{s}, \mathbf{t}) = w_1 \frac{1}{\int_{\Omega} R(x) dx} \int_{\Omega} R(x) P(x, \mathbf{s}) dx - (1 - w_1) \frac{1}{N} \sum_{i=1}^N t_i,$$

where $w_1 \in (0, 1]$ (resp. $1 - w_1$) and $\int_{\Omega} R(x) dx$ (resp. N) are weights associated with the coverage performance metric (resp. cost function). Observing that each component above is properly normalized in $[0, 1]$, we can adopt (4) as long as β is selected so that

$$\beta = \frac{1 - w_1}{w_1} \frac{\int_{\Omega} R(x) dx}{N}. \quad (6)$$

Note that with $t_i \in \{0, 1\}$, the agent heterogeneity in \hat{p}_i (which depends on the values of p_{i0} and λ_i in (1) and on

the sensing range δ_i) is not included in the formulation (4). In order to capture this aspect of the problem, we relax the binary nature of t_i by allowing it to be a continuous variable $t_i \in [0, 1]$. We then rewrite the detection probability in (1) as $t_i p_{i0} e^{-\lambda_i \|x - s_i\|}$ so that t_i acts as a discount factor for the sensing capacity p_{i0} . Accordingly, (2) is modified to

$$\bar{p}_i(x, s_i, t_i) = \begin{cases} t_i p_{i0} e^{-\lambda_i \|x - s_i\|} & \text{if } x \in V(s_i), \\ 0 & \text{otherwise.} \end{cases} \quad (7)$$

and the definition of the joint detection probability in (5) becomes

$$\bar{P}(x, \mathbf{s}, t) = 1 - \prod_{i=1}^N (1 - \bar{p}_i(x, s_i, t_i)).$$

With $\bar{P}(x, \mathbf{s}, t)$ as defined above, we now extend (4) to

$$\begin{aligned} \max_{\mathbf{s}, t} \int_{\Omega} R(x) \bar{P}(x, \mathbf{s}, t) dx - \beta \sum_{i=1}^N t_i \quad (8) \\ \text{s.t. } s_i \in F, \quad t_i \in [0, 1], \quad i = 1, \dots, N. \end{aligned}$$

However, this formulation still does not capture the fact that agents with different sensing parameter values p_{i0} , δ_i and λ_i have different costs. Therefore, let $\gamma_i(p_{i0}, \lambda_i, \delta_i)$ denote the cost of agent i and let us still keep β as a weight indicating the overall relative importance of cost relative to the coverage performance expressed by the first term in the objective function. Omitting the dependence of γ_i on the sensing parameters, we now formulate the problem:

$$\begin{aligned} \max_{\mathbf{s}, t} H(\mathbf{s}, t) = \int_{\Omega} R(x) \bar{P}(x, \mathbf{s}, t) dx - \beta \sum_{i=1}^N \gamma_i t_i \quad (9) \\ \text{s.t. } s_i \in F, \quad t_i \in [0, 1], \quad i = 1, \dots, N. \end{aligned}$$

Clearly, heterogeneity here is captured in two ways: first, by imposing a different cost γ_i to each agent and second by associating a different sensing capacity $t_i p_{i0}$ in (7) to each agent, assuming that such capacity is adjustable. More importantly, while the binary constraint in (4) is removed, the l_1 norm used in (9) is a regularization term which is well known to induce sparsity (e.g., in the use of machine learning algorithms such as LASSO [21]). The implication is that solutions of this problem will tend to include values $t_i = 0$ for several agents in seeking cost-effective team compositions. This is both theoretically proven in Theorem 3, Section 4 and experimentally validated using numerical results in Section 5.

As in the case of (4), the objective function in (9) needs to be properly normalized. To accomplish this while also providing a physical interpretation to the cost coefficients γ_i , recall that $\Omega_i = \{x \mid \|x - s_i\| \leq \delta_i\}$ represents the sensing region of agent i and define the *sensing capability* of this agent as

$$\kappa_i = \int_{\Omega_i} \hat{p}_i(x, s_i) dx,$$

where $s_i \in \mathbb{R}^2$ can be any point in the boundless and obstacle-free space, hence κ_i is independent of s_i ; it depends only on the sensing parameters p_{i0} , λ_i , and δ_i . In fact, for the exponential sensing function given in (1), a closed-form expression for κ_i can be obtained as

$$\kappa_i = \frac{2\pi p_{i0}}{\lambda_i^2} [1 - (1 + \lambda_i \delta_i) e^{-\lambda_i \delta_i}].$$

Now, assuming the cost γ_i associated with agent i is proportional to its sensing capability κ_i , we write:

$$\gamma_i = w_{2i} \kappa_i, \quad (10)$$

where $w_{2i} \in (0, 1]$ is a prespecified *agent cost weight*. Finally, we update the definition of the normalization factor β in (6) as follows:

$$\beta = \frac{1 - w_1}{w_1} \frac{\int_{\Omega} R(x) dx}{\sum_{i=1}^N \gamma_i}. \quad (11)$$

With that, we can compute all the parameters/coefficients behind the formulated optimal agent team composition problem (9), when the agent sensing capabilities and weights (i.e., w_1 and w_{2i} , $\forall i$) are given. Therefore, the problem formulation is now complete.

3 Greedy Algorithm and Submodularity Theory for Coverage Problems

In order to obtain an initial solution to the problem in (9), we first consider the problem given in (3) where the objective is limited to maximizing the coverage using all the available agents. Let us start by adopting the generic greedy method proposed in [20] and seek to improve upon the performance bounds provided in [20] by exploiting the curvature concepts proposed in [8, 14].

3.1 Set-function approach to the basic coverage problem

In order to take advantage of the submodular structure of $H(\mathbf{s})$ in (3), we first uniformly discretize the continuous feasible space F to form a *ground-set* $F^D = \{x_1, x_2, \dots, x_n\}$ with each $x_i \in F$. These x_i values can be thought of as feasible points where an agent can be placed. Note that the cardinality $|F^D|$ of the ground-set is $|F^D| = n$. As the next step, a *set-variable* is defined as $S = \{s_1, s_2, \dots\}$ to represent the initial placement for each agent. Typical constraints on selecting S include the fact that each s_i should be chosen from the ground-set F^D and the total number of agents should be constrained to N . Therefore, the set-constraint $S \in \mathcal{S}$, where $\mathcal{S} = \{A : A \subseteq F^D, |A| \leq N\}$ is used. Typically, a set-constraint of this form is called a *uniform matroid constraint of rank N* where the pair $\mathcal{M} = (F^D, \mathcal{S})$ is known as a *uniform matroid*.

Furthermore, throughout this section, we approximate the coverage objective function $H(\mathbf{s})$ in (3) by a set-function

$H(S)$, where $H : \mathcal{S} \rightarrow \mathbb{R}$ and

$$H(S) = \int_{\Omega} R(x) \left(1 - \prod_{s_i \in S} [1 - \hat{p}_i(x, s_i)]\right) dx. \quad (12)$$

Therefore, $H(S)$ now represents the coverage objective value achieved by the agent placement defined by the set-variable S . In this new framework, a set-function version of the original coverage problem in (3) can be written as

$$\max_S H(S) \quad \text{s.t. } S \in \mathcal{S}. \quad (13)$$

3.2 Greedy algorithm

Due to the combinatorial search space size, an exact solution to (13) is challenging to obtain. However, a candidate solution can be obtained using a simple greedy algorithm and is referred to as a *greedy solution*. Here, we follow the greedy method given in Algorithm 1 to obtain the corresponding greedy solution.

The marginal gain in the coverage objective due to adding a new agent at point $x_i \in F^D$ to an existing agent set A is denoted by $\Delta H(x_i|A)$ where

$$\Delta H(x_i|A) = H(A \cup \{x_i\}) - H(A). \quad (14)$$

This is also known as the *discrete derivative* of the set function $H(S)$ at $S = A$ in the direction x_i [8]. It can be shown through a straightforward evaluation based on (12) that

$$\Delta H(x_i|A) = \int_F R(x) p_i(x, s_i) \prod_{s_j \in A} [1 - \hat{p}_j(x, s_j)] dx. \quad (15)$$

Algorithm 1 Greedy Method for Solving (13)

- 1: **Inputs:** N, F^D (Recall $\mathcal{S} := \{A : A \subseteq F^D, |A| \leq N\}$).
 - 2: **Outputs:** Greedy solution S^G .
 - 3: $S := \emptyset; i := 1;$
 - 4: **while** $i \leq N$ **do**
 - 5: $s_i := \arg \max_{\{x_i : (S \cup \{x_i\}) \in \mathcal{S}\}} (\Delta H(x_i|S));$
 - 6: $S := S \cup \{s_i\};$
 - 7: **end while**
 - 8: $S^G := S;$ **Return;**
-

Using the properties of the problem (13), we can now derive several bounds allowing us to quantify how close the greedy solution is to the globally optimal solution.

3.3 Performance Bounds

Consider the greedy solution of (13) given by Algorithm 1 as $S = S^G$. The *performance ratio* of this greedy solution is defined as $H(S^G)/H(S^*)$ where S^* is the globally optimal

solution (of (13)) and is generally unknown. A *performance bound* L is defined as a theoretically imposed lower bound to the performance ratio. Therefore,

$$L \leq \frac{H(S^G)}{H(S^*)} \leq 1. \quad (16)$$

It was proven in [20] that the set-function $H(S)$ has two important properties: *submodularity* and *monotonicity*. Therefore, following the seminal paper [17], the greedy solution to the coverage problem in (13) is characterized by the performance bound $L = L_C$, where

$$L_C = \left(1 - \left(1 - \frac{1}{N}\right)^N\right). \quad (17)$$

We refer to (17) as the *conventional* performance bound.

3.4 Curvature information

For the class of coverage problems we are considering, it is shown in [20] that tighter performance bounds (i.e., performance bounds which are closer to 1 than L_C) can be obtained using the *curvature* information of the objective function $H(S)$. Typically, any measure of *curvature* of a set function $f(A)$ provides additional information about the nature of its growth when new elements are added to the set-variable A . In other words, *curvature* information characterizes the nature of the monotonicity of $f(A)$. For example, the marginal gain of a coverage objective set-function $H(S)$ (represented by $\Delta H(\cdot|S)$), can drastically drop when elements are added to the set S . Due to this reason, characterizing the set function's monotonicity (using curvature information) can yield vital information about the effectiveness of greedy methods.

3.4.1 Total Curvature

The concept of *total curvature* for generic submodular monotone set-functions was introduced in [8]. When this concept is applied to the class of coverage control problems, the total curvature of $H(S)$ denoted by α_T is given by

$$\alpha_T = \max_{x_i : x_i \in F^D} \left[1 - \frac{\Delta H(x_i|F^D \setminus x_i)}{\Delta H(x_i|\emptyset)}\right], \quad (18)$$

where we use \emptyset to denote the empty set. Further, “ \setminus ” is used to denote the set-subtraction operation (i.e., $A \setminus B = A \cap B^c$). The use of the prefix “total” comes from the fact that α_T is evaluated based on the marginal gain $\Delta H(\cdot|S)$ when $S = \emptyset$ and when $S = F^D \setminus s_j$ (i.e., at extreme ends of possible sets S). Therefore, the total curvature measure tries to characterize the monotonicity of $H(S)$ using its marginal gain evaluated at two extreme ends of choices for S . In the context of real-valued functions defined on a finite interval, the use of total curvature (for monotone submodular set functions) is analogous to attempting to characterize the

shape of a monotonically increasing curve with monotonically decreasing gradient, using only its gradient at its two endpoints.

Using (15), (12), and the knowledge of F^D , the total curvature α_T of the set-function $H(S)$ can be explicitly evaluated. In [8], it is shown that when maximizing a submodular monotone set function with a total curvature α_T , the greedy solution will follow the performance bound $L = L_T$ where

$$L_T = \frac{1}{\alpha_T} \left[1 - \left(\frac{N - \alpha_T}{N} \right)^N \right]. \quad (19)$$

This total curvature measure has been used in [20] to establish better performance bounds compared to the conventional bound L_C in the context of the coverage control problem in (13). Next, we propose another curvature concept to obtain even tighter performance bounds than L_T .

3.4.2 Partial Curvature

In [14], the concept of *partial curvature* is proposed for submodular monotone set functions which are defined under uniform matroid constraints. Adopting this new concept, the partial curvature measure associated with the coverage objective set-function $H(S)$ can be expressed as α_P where

$$\alpha_P = \max_{(A, x_i): x_i \in A \in \mathcal{S}} \left[1 - \frac{\Delta H(x_i | A \setminus x_i)}{\Delta H(x_i | \emptyset)} \right]. \quad (20)$$

As discussed in [14], the partial curvature delivers a better characterization of the monotonicity of any generic set-function compared to the total curvature. This improvement is due to the fact that only the information obtained from the domain of the considered set-function is used - which can be considerably smaller due to the uniform matroid constraint. The importance of the partial curvature concept in the context of our coverage problem can be explained as follows. For coverage problems, evaluating $H(F^D)$ so as to compute the total curvature in (18) and then to impose the performance bound L_T is problematic because the domain of $H(\cdot)$ in the original optimization problem (13) is actually limited to size N sets (i.e., by the constraint $S \in \mathcal{S}$). This issue is critical when we consider heterogeneous agents (in terms of sensing capabilities) and a finite set of agents at our disposal to achieve the maximum coverage. In such situations, $H(F^D)$ is ill-defined and, therefore, the total curvature and the respective performance bound L_T cannot be evaluated. However, the definition of the partial curvature in (20) will still hold as it only requires evaluations of $H(\cdot)$ over the same domain (i.e., $S \in \mathcal{S}$).

Remark 1 *When all the agents available are homogeneous (as opposed to the heterogeneous situation discussed above) the definition of the coverage objective function $H(\cdot)$ in (12) is flexible enough so that its domain can be extended to 2^{F^D}*

(from \mathcal{S}). Thus, it enables the evaluation of the total curvature measure in (18) and the associated performance bound L_T . However, the effectiveness of the bound L_T is questionable since this has been computed using a larger objective function domain (2^{F^D}) while the original optimization problem in (13) is considered over a smaller domain \mathcal{S} . Therefore, it is natural to presume that the total curvature-based performance bound L_T can be further improved when the optimization problem is over a smaller domain.

Using (15), (12) and the knowledge of \mathcal{S} , the partial curvature α_P in (20) can be computed for the coverage problem. The corresponding performance bound is denoted by $L = L_P$, where

$$L_P = \frac{1}{\alpha_P} \left[1 - \left(\frac{N - \alpha_P}{N} \right)^N \right]. \quad (21)$$

3.4.3 Greedy Curvature

We also introduce the use of another curvature concept, the *greedy curvature*, which is proposed in [8] as an on-line method of estimating a performance bound. The resulting performance bound depends on the greedy solution S^G itself. Note that the performance bounds discussed thus far are not dependent on the obtained greedy solution but only the objective function parameters (such as λ_i, δ_i for all i) and N , as well as the feasible space F^D .

If the greedy algorithm given in Algorithm 1 produces the solution sets $\emptyset = S^0 \subseteq S^1 \subseteq S^2 \subseteq \dots \subseteq S^N$ during the course of execution (where $S^N = S^G$), then, the greedy curvature metric α_G is given by

$$\alpha_G = \max_{0 \leq i \leq N-1} \left[\max_{x_j \in F^i} \left(1 - \frac{\Delta H(x_j | S^i)}{\Delta H(x_j | \emptyset)} \right) \right], \quad (22)$$

where $F^i = \{x_j : x_j \in F^D \setminus S^i, (S^i \cup \{x_j\}) \in \mathcal{S}\}$ is the set of valid points considered for the placement of the $(i+1)^{\text{th}}$ agent during the $(i+1)^{\text{th}}$ greedy iteration. Therefore, α_G can be computed in parallel with the greedy method (without performing any additional computations) unlike the previously discussed two cases. The corresponding performance bound denoted by $L = L_G$ is

$$L_G = 1 - \alpha_G \left(1 - \frac{1}{N} \right). \quad (23)$$

The main idea behind the greedy curvature concept is that the solution sets generated during the greedy algorithm itself can be used to characterize the monotonicity of the considered set-function and then to establish a performance bound based on that information. Therefore, similar to the observation made earlier regarding the feasibility of using the total curvature-based performance bound L_T for a heterogeneous

set of agents, the definition of the greedy curvature measure in (22) and the performance bound L_G in (23) will still hold in such cases.

3.5 The Overall Performance Bound L

Taking all the aforementioned performance bounds L_C, L_T, L_P and L_G defined respectively in (17), (19), (21), and (23), into account, an overall performance bound L satisfying (16) can be established as

$$L = \max \{L_C, L_T, L_P, L_G\}. \quad (24)$$

Generally, $L_C \leq L_T \leq L_P$ [14]. Also, recall that when heterogeneous agents are involved, L_T and L_G are undefined.

3.6 Numerical results for greedy method

We now investigate the behavior of the proposed partial curvature and greedy curvature-based performance bounds L_P, L_G compared to the conventional and total curvature performance bounds L_C, L_T . Four different representative problem settings were considered as shown in Fig. 2. Under each of these settings, the aforementioned performance bounds were evaluated for different values of the total allowable number of agents N .

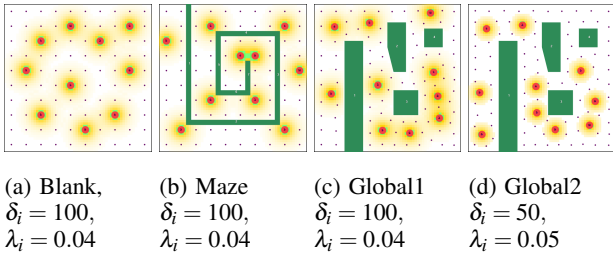


Fig. 2. Different problem settings and their greedy solutions for $N = 10$. Red dots are greedy agent locations, black dots represent the ground set. Darker colored areas have greater coverage, and green colored shapes are obstacles.

From the obtained results shown in Fig. 3, it is evident that the proposed use of partial curvature always delivers better bounds than the total curvature approach [8]. Similarly, the proposed use of greedy curvature provides better bounds than the total curvature approach [8] when N takes moderate values (i.e., N is around 2 – 20). Moreover, L_G is useful for computation-limited settings, as it does not require any additional computations compared to evaluating L_T or L_P .

As pointed out earlier, the performance bound L_T is ill-defined when considering heterogeneous agents. To avoid this problem, the experiments reported above were limited to a homogeneous set of agents. However, it should be emphasized that the definitions of the proposed performance bounds L_P and L_G are robust to agent heterogeneity, the situation considered in section 5. Therefore, in such heterogeneous situations, using L_P and/or L_G will be the only




Fig. 3. Performance bounds (as a function of N): (i) Conventional L_C , (ii) Total curvature L_T , (iii) Partial curvature L_P , and (iv) Greedy curvature L_G , for the four problem settings in Fig. 2.

way to obtain an improved performance bound compared to the conventional bound L_C . Note that in such situations, the greedy algorithm given will require an additional inner loop to determine the optimal type of the agent to be deployed at each main greedy iteration.

We conclude this section by reminding the reader that the greedy process detailed above is needed so as to generate an initial condition to the main optimization problem in (9). We have also discussed different performance bound computation techniques which can characterize the closeness of these initial conditions to the global optimum.

4 A Gradient Based Algorithm for Heterogeneous Multi-Agent Coverage Problem

The greedy algorithm (Algorithm 1) is limited to discrete environments and a fixed predetermined agent number. Its value in solving the actual problem of interest in (9) is twofold: (i) Provide a reasonable initial condition for a gradient-based algorithm used to solve (9) which can significantly overcome the local-optimality limitation of such an algorithm, and (ii) Provide a lower bound for the ultimate coverage performance we obtain.

In this section, we propose a distributed gradient-based algorithm similar to that in [28] aimed at solving (9). We first derive the derivatives of the objective function $H(s, t)$ with regard to the variables (s, t) for the gradient ascent update. Setting $s_i = (s_{ix}, s_{iy})$, we begin with $\frac{\partial H(s, t)}{\partial s_{ix}}$ whose derivation

was given in [28]:

$$\begin{aligned} \frac{\partial H(\mathbf{s}, t)}{\partial s_{ix}} &= \int_{V(s_i)} R(x) \Phi_i(x) \frac{\partial \bar{p}_i(x, s_i, t_i)}{\partial s_{ix}} dx \\ &+ \sum_{j \in \Gamma_i} \text{sgn}(n_{ijx}) \frac{\sin(\theta_{ij})}{D_{ij}} \int_0^{Z_{ij}} R(\rho(r)) \Phi_i(\rho(r)) \bar{p}_i(\rho(r), s_i, t_i) r dr, \end{aligned} \quad (25)$$

where

$$\begin{aligned} \Phi_i(x) &= \prod_{k \in B_i} [1 - \bar{p}_k(x, s_k, t_k)], \\ \frac{\partial \bar{p}_i(x, s_i, t_i)}{\partial s_{ix}} &= -\lambda_i \bar{p}_i(x, s_i, t_i) \left(\frac{(s_i - x)_x}{\|s_i - x\|} \right), \\ \rho(r) &= \rho_{ij}(r) = \left(\frac{v_{ij} - s_i}{D_{ij}} \right) r + v_{ij}, \quad \text{and,} \\ D_{ij} &= \|v_{ij} - s_i\|. \end{aligned}$$

In (25), $\text{sgn}(\cdot)$ represents the signum function and the subscript x is used to represent the x -component of a two dimensional vector. The second term in (25) is due to the linear shaped boundary segments of the sensing region $V(s_i)$ formed due to the obstacle vertices $v_{ij} \in V(s_i)$. Such linear segments are lumped into a set $\Gamma_i = \{\Gamma_{i1}, \Gamma_{i2}, \dots\}$ where each linear segment Γ_{ij} can be characterized by four parameters: (i) end point Z_{ij} , (ii) angle θ_{ij} , (iii) obstacle vertex v_{ij} , and, (iv) unit normal direction n_{ij} . Therefore, Γ_{ij} can be thought of as a four-tuple $\Gamma_{ij} = (Z_{ij}, \theta_{ij}, v_{ij}, n_{ij})$. All these geometric parameters (for a generic setting) are illustrated in Fig. 2. Note that we assume: (i) obstacles are polygonal, and, (ii) sensing power at the edge of the sensing region is negligible. More detailed definitions and derivations are omitted for brevity, and interested readers are referred to [28].

A similar expression can be obtained for $\frac{\partial H(\mathbf{s}, t)}{\partial s_{iy}}$. As detailed in [28], the agent locations are assumed not to coincide with a reflex vertex, a polygonal inflection, or a bi-tangent where $H(\mathbf{s}, t)$ is not differentiable (if such points have to be taken into consideration, then a subgradient can be used as an alternative to the gradient).

Additionally, the derivative $\frac{\partial H(\mathbf{s}, t)}{\partial t_i}$ is obtained as follows:

$$\frac{\partial H(\mathbf{s}, t)}{\partial t_i} = \underbrace{\int_{V(s_i)} R(x) \Phi_i(x) p_i(x, s_i) dx}_{\text{Local Coverage}} - \underbrace{\beta \gamma_i}_{\text{Local Cost}}. \quad (26)$$

Here, the integration and differentiation are interchangeable since $P(x, \mathbf{s})$ is a continuous differentiable function of t_i . The first term in (26) represents a *local coverage* level achieved by the agent i in its sensing region $V(s_i)$. This local coverage level depends on the state variables (\mathbf{s}, t) and is always positive. The second term in (26) represents a *local cost* resulting from agent cost γ_i and the normalization factor β .

Note that this local cost value is a predefined positive constant for each agent. This multi-objective interpretation of (26) can be used to conclude that when the aforementioned local coverage level is less than the (fixed) local cost, the state variable t_i should be decreased to improve the global objective $H(\mathbf{s}, t)$, and vice versa.

Algorithm 2 is a Projected Gradient Ascent (PGA) algorithm for solving (9) which utilizes the gradients derived in (25) and (26). As seen in Algorithm 2, a gradient ascent update is first implemented in (29), where $\eta_s^{(k)} > 0$, $\eta_t^{(k)} > 0$ are the step sizes chosen based on standard technical conditions [2] (more application-specific details on the step size selection can be found in [24]). Subsequently, the projection mechanisms are applied to guarantee the satisfaction of all constraints. The projection $\Pi_A(x)$ of $x \in \mathbb{R}^n$ onto a set $A \subseteq \mathbb{R}^n$ is formally defined as

$$\Pi_A(x) \triangleq \arg \min_{y \in A} \|x - y\|_2. \quad (27)$$

For $s_i \in F$, if the update direction (i.e., $\frac{\partial H(\mathbf{s}, t)}{\partial s_i}$) is pointing directly into an obstacle's boundaries, then the update direction is projected onto the boundary itself and thus prevents violation of the obstacle constraint. As for the bound constraint for t_i , a projection onto the convex set $[0, 1]$ is simply a truncation.

Coverage performance of the PGA solution \mathbf{s}^{PGA} For the initialization of the PGA algorithm given in Algorithm 2, we use the greedy solution $\mathbf{s}^{(0)} = [S^G]$ obtained from Algorithm 1 using: (i) The pre-specified discretized feasible space F^D , and, (ii) The complete set of agents (all N of them). The overall performance bound obtained using (24) under this initial configuration is $L_1 \leq \frac{H(\{\mathbf{s}^{(0)}\})}{H(S^*)}$. Therefore, L_1 does not convey any information about the coverage performance of the obtained PGA solution. This issue is addressed as follows (using the notation $[\cdot]$ and $\{\cdot\}$ to represent a conversion from a set to an array and vice versa).

Once the PGA solution $(\mathbf{s}^{PGA}, t^{PGA})$ is obtained using Algorithm 2, it yields information on: (i) optimal agent locations (i.e., \mathbf{s}^{PGA}), and, (ii) optimal agent team composition (i.e., t^{PGA}). This allows us to update the discretized feasible space F^D into F^{D2} by inserting the agent coordinates found in \mathbf{s}^{PGA} such that $F^{D2} \triangleq F^D \cup \{\mathbf{s}^{PGA}\}$. Next, we re-evaluate the greedy algorithm considering only the agents in the optimal team and using the modified discretized feasible space F^{D2} . Now, if the corresponding greedy solution is S^{G2} and the overall performance bound is L_2 (obtained from (24)), following (16) we can write $L_2 \leq \frac{H(S^{G2})}{H(S^*)}$. This relationship together with $H(\{\mathbf{s}^{PGA}\})$ can then be used to impose a lower

Algorithm 2 Projected Gradient Ascent (PGA) Algorithm for solving the problem in (9).

- 1: **Inputs:** Ω , F , N , and, tolerances $\varepsilon_s, \varepsilon_t > 0$.
- 2: **Initialize:** $\mathbf{s}^0 := [S^G]$ (From Alg. 1), $\mathbf{t}^0 := [1, 1, \dots, 1] \in \mathbb{R}^n$.
- 3: **Outputs:** $\mathbf{s}^{PGA}, \mathbf{t}^{PGA}$.
- 4: **for** $k = 0, 1, 2, \dots$ **do**
- 5: **Compute:** At $(\mathbf{s}, \mathbf{t}) = (\mathbf{s}^{(k)}, \mathbf{t}^{(k)})$, $\forall i \in \{1, 2, \dots, N\}$,

$$\frac{\partial H(\mathbf{s}, \mathbf{t})}{\partial s_i} = \left[\frac{\partial H(\mathbf{s}, \mathbf{t})}{\partial s_{ix}}, \frac{\partial H(\mathbf{s}, \mathbf{t})}{\partial s_{iy}} \right]; \quad (28)$$

$$\frac{\partial H(\mathbf{s}, \mathbf{t})}{\partial t_i};$$

▷ Using (25) and (26).

- 6: **Update:** $(\mathbf{s}^{(k)}, \mathbf{t}^{(k)})$, by, $\forall i \in \{1, 2, \dots, N\}$,

$$\hat{s}_i^{(k+1)} = s_i^{(k)} + \eta_s^{(k)} \frac{\partial H(\mathbf{s}, \mathbf{t})}{\partial s_i}; \quad (29)$$

$$\hat{t}_i^{(k+1)} = t_i^{(k)} + \eta_t^{(k)} \frac{\partial H(\mathbf{s}, \mathbf{t})}{\partial t_i};$$

- 7: **Projection:** to get $(\mathbf{s}^{(k+1)}, \mathbf{t}^{(k+1)})$, $\forall i \in \{1, 2, \dots, N\}$,

$$s_i^{(k+1)} = \Pi_F(\hat{s}_i^{(k+1)}); \quad (30)$$

$$t_i^{(k+1)} = \Pi_{[0,1]}(\hat{t}_i^{(k+1)});$$

- 8: **if** $\{\|\mathbf{s}^{(k+1)} - \mathbf{s}^{(k)}\| \leq \varepsilon_s$ and $\|\mathbf{t}^{(k+1)} - \mathbf{t}^{(k)}\|\} \leq \varepsilon_t$ **then**
 - 9: $(\mathbf{s}^{PGA}, \mathbf{t}^{PGA}) := (\mathbf{s}^{(k+1)}, \mathbf{t}^{(k+1)})$; **Return;**
 - 10: **end if**
 - 11: **end for**
-

bound to the ratio $\frac{H(\{\mathbf{s}^{PGA}\})}{H(S^*)}$ as follows:

$$L' \triangleq L_2 \cdot \frac{H(\{\mathbf{s}^{PGA}\})}{H(S^{G2})} \leq \frac{H(\{\mathbf{s}^{PGA}\})}{H(S^*)}. \quad (31)$$

Therefore, L' can be used as a performance bound guarantee on the final coverage level achieved by the chosen optimal team of agents.

Characterization of optimal t_i values given by PGA We consider two agents i and j to be *neighbors* if their sensing regions overlap (i.e., $V(s_i) \cap V(s_j) \neq \emptyset$). The *set of neighbors* of agent i is denoted by $B_i = \{j : j \neq i, V(s_i) \cap V(s_j) \neq \emptyset\}$. Note that B_i does not include i , therefore, we define the *closed neighborhood* of agent i as $\bar{B}_i = B_i \cup \{i\}$. Using these neighborhood concepts, we define the following state variable compositions to go along with (s_i, t_i) :

- The *neighbor state* variables: $(\bar{s}_i^c, \bar{t}_i^c)$, where $\bar{s}_i^c = [\{s_j : j \in B_i\}]$ and $\bar{t}_i^c = [\{t_j : j \in B_i\}]$.
- The *neighborhood state* variables: (\bar{s}_i, \bar{t}_i) , where $\bar{s}_i = [\{s_j : j \in \bar{B}_i\}]$ and $\bar{t}_i = [\{t_j : j \in \bar{B}_i\}]$.
- The *complementary state* variables: (s_i^c, t_i^c) , where $s_i^c = [\{s_j : \forall j \neq i\}]$ and $t_i^c = [\{t_j : \forall j \neq i\}]$.

Using this notation, we can now establish the following lemma.

Lemma 2 The objective function $H(\mathbf{s}, \mathbf{t})$ in (9) can be decomposed as,

$$H(\mathbf{s}, \mathbf{t}) = t_i H_i(\bar{s}_i, \bar{t}_i^c) + H_i^c(s_i^c, t_i^c) \quad (32)$$

where

$$H_i(\bar{s}_i, \bar{t}_i^c) = \int_{V(s_i)} R(x) \Phi_i(x) p_i(x, s_i) dx - \beta \gamma_i, \text{ and,}$$

$$H_i^c(s_i^c, t_i^c) = \int_{\Omega} R(x) \left[1 - \prod_{\forall l \neq i} (1 - \bar{p}_l(x, s_l, t_l)) \right] dx - \beta \sum_{\forall l \neq i} \gamma_l t_l.$$

Proof: $H(\mathbf{s}, \mathbf{t})$ as given in (9) can be expanded as

$$H(\mathbf{s}, \mathbf{t}) = \int_{\Omega} R(x) (1 - (1 - \bar{p}_i(x, s_i, t_i)) \prod_{\forall l \neq i} (1 - \bar{p}_l(x, s_l, t_l))) dx - \beta \gamma_i t_i - \beta \sum_{\forall l \neq i} \gamma_l t_l. \quad (33)$$

Now, using the following relationships directly obtained from (7), (2), along with the definition of the neighbor set B_i :

$$\begin{aligned} \bar{p}_i(x, s_i, t_i) &= t_i p_i(x, s_i) & \forall x, s_i \in \Omega, \forall t_i \in [0, 1], \\ p_i(x, s_i) &= 0 & \forall s_i \in \Omega, x \notin V(s_i), \\ p_i(x, s_i) (1 - p_j(x, s_j)) &= p_i(x, s_i) & \forall x \in \Omega, \forall j, j \notin B_i, \end{aligned}$$

we can write, for all $x, s_i, s_l \in \Omega$ and $t_i, t_l \in [0, 1]$,

$$\bar{p}_i(x, s_i, t_i) \prod_{\forall l \neq i} (1 - \bar{p}_l(x, s_l, t_l)) = t_i p_i(x, s_i) \prod_{l \in B_i} (1 - \bar{p}_l(x, s_l, t_l)).$$

Using the above relationship in (33), we obtain (32). ■

Using Lemma 2 we establish the following theorem which characterizes the nature of t_i^* , the t_i values given by the PGA Algorithm 2.

Theorem 3 For any agent i , the values obtained from the PGA algorithm satisfy

$$t_i^* = \begin{cases} 0 & \text{when } H_i(\bar{s}_i^*, \bar{t}_i^{c*}) < 0, \\ 1 & \text{when } H_i(\bar{s}_i^*, \bar{t}_i^{c*}) > 0. \end{cases} \quad (34)$$

Moreover, when $H_i(\bar{s}_i^*, \bar{t}_i^{c*}) = 0$, the optimal objective function value $H(\mathbf{s}^*, \mathbf{t}^*)$ is invariant to t_i^* .

Proof: Using the decomposition shown in Lemma 2, we get

$$\frac{\partial H(\mathbf{s}, t)}{\partial t_i} = H_i(\bar{s}_i, \bar{t}_i^c),$$

where $H_i(\bar{s}_i, \bar{t}_i^c)$ is independent of t_i . Therefore, when $H_i(\bar{s}_i, \bar{t}_i^c) \neq 0$, it is clear that the PGA cannot terminate the t_i update process in (29) until t_i hits a constraint boundary given by $t_i \in [0, 1]$. The update direction depends on the sign of $H_i(\bar{s}_i, \bar{t}_i^c)$ and update process in (29) will become stationary when t_i satisfies (34).

To prove the second statement, consider the case where $H_i(\bar{s}_i^*, \bar{t}_i^{c*}) = 0$ with $t_i^* \in (0, 1)$. Since $H_i(\bar{s}_i, \bar{t}_i^c)$ is independent of t_i , if t_i^* is perturbed to a value $t_i = t_i^* + \Delta \in [0, 1]$, the optimality condition $H_i(\bar{s}_i^*, \bar{t}_i^{c*}) = 0$ still holds true. Further, using this relationship with Lemma 2, we can see that $H(\mathbf{s}, t)$ is insensitive to a perturbation $t_i^* + \Delta \in [0, 1]$ when at $(\mathbf{s}, t) = (\mathbf{s}^*, t^*)$. This means that if the PGA converges to a value $t_i = t_i^* \in (0, 1)$, perturbing t_i towards either 0 or 1 will not affect the objective function value. This concludes the proof. \blacksquare

Remark 4 *Using Lemma 2, it can be further shown that, when $H_i(\bar{s}_i^*, \bar{t}_i^{c*}) = 0$, with $t_i^* \in (0, 1)$, if t_i^* is artificially perturbed, the optimality condition for s_i (i.e., $\frac{\partial H(\mathbf{s}, t)}{\partial s_i} = 0$) still holds. However, due to such a perturbation, the optimality conditions of neighbor agent states are affected (i.e., $\frac{\partial H(\mathbf{s}, t)}{\partial s_j} \neq 0$, $\frac{\partial H(\mathbf{s}, t)}{\partial t_j} \neq 0$, $j \in B_i$). In a such situation, the PGA should be re-activated from the perturbed state. Also note that in numerical simulations, occurrence of a such equivalence is unlikely.*

In conclusion, the proposed PGA ensures that the resulting optimal t_i values are either 0 or 1. Hence, despite the relaxation of the binary variable t_i to $t_i \in [0, 1]$, it provides a solution to the mixed integer non-linear programming problem version of (9), where, for all i , t_i is constrained to $t_i \in \{0, 1\}$.

We conclude this section by observing that Lemma 2 makes it clear that in order for an agent to compute the gradients required in (28) (i.e., at step 5 of Algorithm 2), it only needs the neighborhood state information (\bar{s}_i, \bar{t}_i) . Therefore, in executing the PGA, agents have the capability to perform all required computations (and subsequent actuations) in a distributed manner.

5 Numerical Results

In this section, we provide several numerical results obtained from the proposed PGA (Algorithm 2) initialized with the solution provided by the greedy Algorithm 1 discussed in Section 3. The PGA method is evaluated under four different mission space configurations named: (i) General, (ii) Room, (iii) Maze, and, (iv) Narrow, as shown in Figs. 5a, 5b, 5c and 5d, respectively. The mission space is a square of size 600×600 units with an event density function $R(x)$

assumed to be uniform (i.e., $R(x) = 1$, $\forall x \in F$). All simulations are initialized with ten agents (i.e., $N = 10$) and each agent's nominal sensing capacity is selected as $p_{i0} = 1$. For the use of the greedy algorithm, the ground set F^D is constructed by uniformly placing 100 points in the mission space. All reported simulation results and execution times have been obtained by executing the algorithms on a standard desktop computer with 8.0 GB RAM and a 3.61 GHz AMD eight-core processor. For convenience, we define the cost component of the overall objective function $H(\mathbf{s}, t)$ as $C(t) = \beta \sum_{i=1}^N \gamma t_i$. Therefore, $H(\mathbf{s}, t) = H(\mathbf{s}) - C(t)$ where $H(\mathbf{s})$ represents the coverage component of $H(\mathbf{s}, t)$.

5.1 The homogeneous agent case

We first consider the case of homogeneous agents with a sensing decay $\lambda_i = 0.012$, and a sensing range $\delta_i = 200$ units in the mission space. The weight parameters are selected as $w_1 = 0.68$ and $w_{2i} = 1, \forall i = 1, \dots, N$ in (9). The obtained results are summarized in Tab. 1 where each row corresponds to one of the four mission space configurations defined above. The first part of the table gives the results of the initial greedy algorithm where the cost component is ignored. The second part gives the final results of the PGA. Figures 5a, 5b, 5c and 5d compare the resulting system configurations at the aforementioned two stages of the PGA. Note that the agents drawn as light-colored disks are those with $t_i = 0$, and, therefore, are not included in the optimal agent team.

The overall objective value improvement over the initial greedy solution can be seen by comparing the two $H(\mathbf{s}, t)$ columns in Tab. 1. This improvement is a result of excluding some of the agents (from $N = 10$) and obtaining solutions with optimal agent team size $N < 10$. It was observed that such agent exclusions (i.e., $t_i = 0$) occur when an agent's terminal location s_i^{PGA} is in: (i) A confined/narrow region where it cannot fully utilize its sensing capabilities (e.g., see agent 10 in Fig 5b), or in, (ii) A region which is already covered by other agents (e.g., see agent 3 in Fig 5a). As expected (see Theorem 3), all observed optimal t_i values converged to either 0 or 1 and without the need of the extra PGA step described in Remark 4.

Moreover, the coverage performance bounds L' (defined in (31)) achieved by the optimal agent teams are listed in Tab. 2. From these results, we can conclude that: (i) On average, the optimal team provides more than 80% of the attainable maximum coverage level, and, (ii) In some mission spaces, we can even guarantee near global optimality (e.g., in the Narrow mission space).

The effect of the normalization factor β We have studied the effect of the normalization factor β , which captures the trade-off between the coverage performance and the team cost. Thus, decreasing the value of β value highlights the effect of coverage $H(\mathbf{s})$ over the team cost $C(t)$ in the overall objective function $H(\mathbf{s}, t)$. Since, from (11), β directly depends on the normalization weight w_1 , we tune w_1 to get

Table 1
Results of the proposed PGA for the homogeneous agent case.

Mission Space	Initial Greedy Solution					Final PGA Solution					Fig.
	N	$H(s)$	$C(t)$	$H(s,t)$	time/s	N	$H(s)$	$C(t)$	$H(s,t)$	time/s	
General	10	157,111	142,289	14,822	2.135	7	127,225	99,645	27,580	1.763	5a
Room	10	145,206	142,289	2,917	2.056	5	94,441	71,215	23,225	2.919	5b
Maze	10	148,082	142,289	5,793	1.888	7	112,915	99,645	13,270	3.112	5c
Narrow	10	184,076	142,289	41,787	2.197	7	150,074	99,645	50,429	2.663	5d

Table 2
Performance bound guarantees (i.e., L' in (31)) on the final coverage level achieved by the optimal agent team for the homogeneous agent case.

Mission Space	N	$H(S^{G^2})$	L_2	L'
General	7	114,804	0.651	0.721
Room	5	93,086	0.874	0.886
Maze	7	112,508	0.665	0.667
Narrow	7	148,073	0.999	0.999

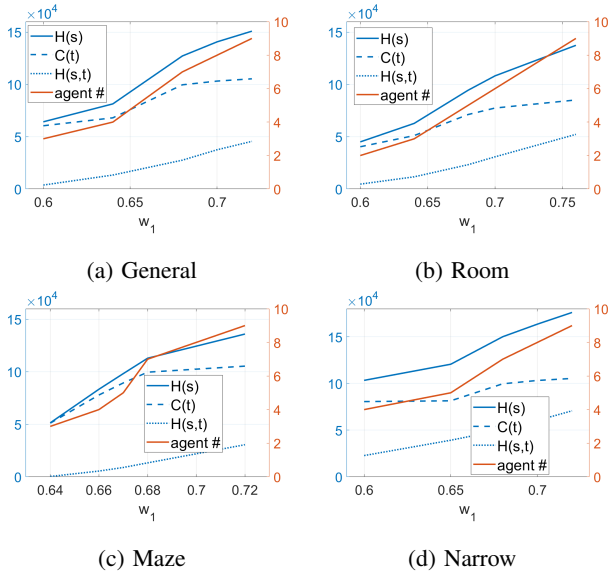




Fig. 4. Effect of the normalization weight w_1 on the obtained PGA solution: $H(s), C(t), H(s,t)$ and N , in different mission spaces for the homogeneous agent case.


different β values while keeping $w_{2i} = 1, \forall i$. Figures 4a, 4b, 4c and 4d show the effect of the normalization weight w_1 on the obtained results (i.e., the achieved $H(s), C(t), H(s,t)$ and N values by the PGA algorithm) for the four different mission space configurations. The main conclusions on the behavior of the observed parameters w.r.t. w_1 are: (i) It generally depends on the considered mission space, (ii) It is non-decreasing and piece-wise linear, and, (iii) $H(s)$ grows faster than $C(t)$.




(a) General



(b) Room



(c) Maze



(d) Narrow

Fig. 5. Comparison of initial greedy solution (left) and projected gradient ascent (PGA) algorithm solution (right) under different mission spaces for the homogeneous agent case.

Table 3
Results of the proposed PGA for the heterogeneous agent case.

Mission Space	Initial Greedy Solution				Final PGA Solution				Fig.
	$N = 10$	$H(s)$	$C(t)$	$H(s,t)$	Agent Team	$H(s)$	$C(t)$	$H(s,t)$	
General	10	152,272	177,140	-22,868	$\{1, 2, 4, 5\}, \{6, 7, 8\}$	124,194	128,174	-3,980	6a
Room	10	142,859	177,140	-34,281	$\{1, 2, 3\}, \{7, 10\}$	94,417	92,781	1,635	6b
Maze	10	146,175	177,140	-30,965	$\{1, 2, 4\}, \{6, 7, 10\}$	96,889	106,355	-9,465	6c
Narrow	10	179,478	177,140	2,337	$\{1, 2, 3, 4, 5\}, \{6, 7\}$	145,963	136,420	9,543	6d

5.2 The heterogeneous agent case


We now consider the heterogeneous agent case where agents differ from each other in terms of both sensing parameters (i.e., sensing range δ_i and sensing decay λ_i) and cost parameters (i.e., agent cost γ_i). To create such a heterogeneous agent configuration, we first assume that the initially available 10 agents belong to two classes (5 agents per each class) as given in Tab. 4. Then, we set the agent cost weights to $w_{2i} = 1 \forall i$. Based on (10), under each adopted agent class, sensing parameters δ_i and λ_i will determine the agent cost γ_i values as shown in Tab. 4 under the ‘‘Case 5.2’’ column. The normalization weight used is $w_1 = 0.58$. To make the problem meaningful, the agent classes have been chosen so that they have complementary sensing properties. For comparison purposes, note that in the previously discussed homogeneous agent case, all agents belonged to Class 1.

The results obtained from the PGA algorithm are summarized in Tab. 3 and the corresponding optimal agent team deployments are shown in Fig. 6, both at the initial greedy step and at the final PGA solution. Similar to the previously discussed homogeneous agent case, we can see the significant improvement achieved in $H(s,t)$ by the PGA steps compared to the initial greedy solution. It is noteworthy that the PGA algorithm has chosen agents from both classes to form the optimal agent team. Also note that, with the help of the initial greedy step, the PGA method has been capable of placing agents in appropriate mission space regions well suited for their specific sensing properties (see agent 6 in Fig. 6a).


The coverage performance bounds L' (defined in (31)) achieved by the optimal agent teams are shown in Tab. 6. From those results, we can conclude that, on average, the optimal agent team provides more than 75% of the attainable maximum coverage level (slightly less than the average bound observed for the homogeneous agent case).

Table 4
Different classes of agents.


Class	Index i	Sensing Para.		Case 5.2	Case 5.3
		Range (δ_i)	Decay (λ_i)	$w_{2i} = 1,$ Cost (γ_i)	$\gamma_i = 30175,$ Weight (w_{2i})
1	1 ~ 5	200	0.012	30175	1.000
2	6 ~ 10	100	0.008	18772	1.607




(a) General



(b) Room



(c) Maze



(d) Narrow

Fig. 6. Comparison of initial greedy solution (left) and projected gradient ascent (PGA) algorithm solution (right) under different mission spaces for the heterogeneous agent case.

Table 5
Results of the proposed PGA for the sensing-wise heterogeneous agent case.

Mission Space	N	Initial Greedy Solution			Final PGA Solution				Fig.
		$H(s)$	$C(t)$	$H(s,t)$	Agent Team	$H(s)$	$C(t)$	$H(s,t)$	
General	10	156,142	177,140	-20,997	$\{1, 2, 3, 4, 5\}, \{\}$	97,398	88,671	8,726	7a
Room	10	145,848	177,140	-31,292	$\{1, 2, 3, 5\}, \{\}$	79,771	70,972	8,798	7b
Maze	10	146,175	177,140	-30,975	$\{1, 2, 3, 4, 5\}, \{\}$	83,261	88,671	-5,410	7c
Narrow	10	179,478	177,140	2,337	$\{1, 2, 3, 4, 5\}, \{\}$	120,374	88,671	31,703	7d

5.3 Sensing-wise heterogeneous agent case

Our purpose here is to highlight the importance of having different agent costs γ_i when the sensing parameters of the agents are different. We also highlight the importance of using the sensing capability (i.e., κ_i) dependent agent costs as proposed in (10). Unlike the previously discussed heterogeneous agent case, here we use a fixed agent cost $\gamma_i = 30175$ across all agent classes. To achieve this under (10), we manipulate the agent cost weight w_{2i} parameters in each agent class, as given in Tab. 4 column ‘‘Case 5.3’’. As a result of this manipulation, despite the differences in sensing parameters over different agents, the agent costs γ_i across all agents become identical. The normalization weight used is $w_1 = 0.58$.

Since all the other problem settings are identical to the previously discussed heterogeneous agent case (in subsection 5.2), the initial greedy step of the PGA algorithm will yield the same agent deployment. However, the associated total agent cost $C(t)$ will be different due to the modification of agent cost parameters w_{2i} and γ_i compared to that of the heterogeneous agent case. The numerical results obtained are summarized in Tab. 5 and the optimal agent team deployments are shown in Fig. 7. The coverage performance bounds L' (defined in (31)) achieved by the optimal agent team are tabulated in Tab. 7.

As expected, when identical agent costs are used despite their differences in sensing capabilities, the resulting PGA solution gives preference to agents with higher sensing capabilities. As a result, the optimal agent team is inherently biased towards Class 1 agents (see Tab. 5 and notice $\kappa_1 > \kappa_2$ due to the δ_i, λ_i values $i = 1, 2$). Clearly, in real-world applications one expects more capable sensors to have higher

Table 6
Performance bound guarantees (i.e., L' in (31)) on the final coverage level achieved by the optimal agent team for the heterogeneous agent case.

Mission Space	Agent Team	$H(S^{G^2})$	L_2	L'
General	$\{2, 3, 4, 5\}, \{6, 9, 10\}$	117,923	0.703	0.740
Room	$\{1, 2, 3\}, \{7, 10\}$	86,534	0.853	0.931
Maze	$\{1, 2, 4\}, \{6, 7, 10\}$	91,203	0.703	0.747
Narrow	$\{1, 2, 3, 4, 5\}, \{6, 7\}$	144,852	0.651	0.656

Table 7
Performance bound guarantees (i.e., L' in (31)) on the final coverage level achieved by the optimal agent team for the sensing-wise heterogeneous agent case.

Mission Space	Agent Team	$H(S^{G^2})$	L_2	L'
General	$\{1, 2, 3, 4, 5\}, \{\}$	95,633	0.729	0.742
Room	$\{1, 2, 3, 5\}, \{\}$	73,864	0.813	0.878
Maze	$\{1, 2, 3, 4, 5\}, \{\}$	82,957	0.703	0.706
Narrow	$\{1, 2, 3, 4, 5\}, \{\}$	117,231	0.651	0.668

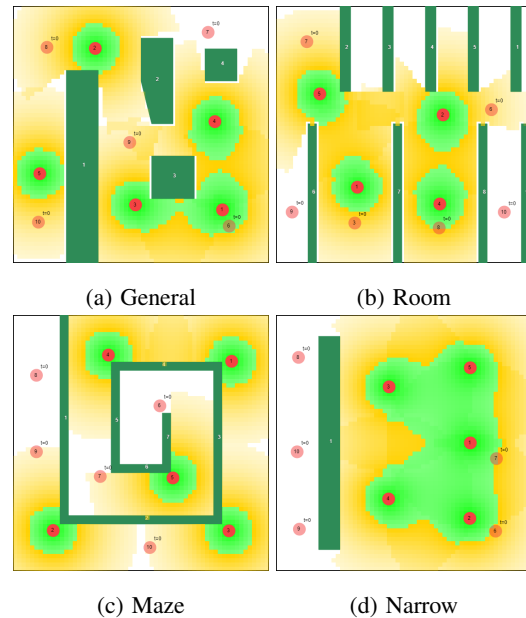


Fig. 7. The obtained final PGA solution under different mission spaces for the sensing-wise heterogeneous agent case.

costs.

5.4 Comparison with a commercial optimization solver

In comparing the solutions given by the proposed PGA method to those of a commercially available optimization problem solver, there are two constraining factors to consider: (i) The coverage component of the objective function in (9) is non-convex, non-linear, and discontinuous. As a result, even though the original version of (9) is a mixed-integer non-linear program (MINLP) (where $t_i \in \{0, 1\}, \forall i$), we were constrained to using a generic non-linear program

(NLP) solver. Therefore, in order to find the optimal binary decision variables (i.e., t), we applied the NLP solver exhaustively over all possible integer values (we refer to this as the “brute force” method). (ii) When obstacles are present in the mission space, the feasible space for each agent becomes non-convex (in our case, this complicates the objective function as well). Since representing such constraints and feeding them to a generic optimization problem solver is difficult, we confine our study to an obstacle-less (blank) mission space.

The NLP solver used is the *interior point method* implemented under the *fmincon* command in MATLAB®. The available agents and their sensing capabilities are given in Tab. 4. In the brute force approach, each iteration considers a specific agent team and computes the optimal coverage solution. Two brute force methods (BF₁ and BF₂) were used depending on the agent initialization in order to highlight the effect of such initialization. Specifically, in BF₁, agents are initialized randomly and in BF₂, agents are initialized in a corner of the mission space such that the l^{th} agent ($\forall l$) is placed at $s_l = (5 + 5l, 5 + 5l)$. Note that when the normalization weight is $w_1 = 1$ (see (11), (9)), the PGA method basically solves the optimal coverage problem. This enables a direct comparison of the performance of the PGA method (when $w_1 = 1$) with that of single iterations of BF₁ and BF₂. This comparison is shown in Fig. 8 and it confirms that the proposed PGA method: (i) Delivers better coverage levels, and, (ii) Shows extremely low execution times compared to BF₁ or BF₂. Another conclusion is that the random initialization has helped the BF₁ method to achieve better coverage and execution times compared to that of BF₂.


Under the information in Tab 4, there are 35 possible agent team configurations. Therefore, 35 brute force iterations were required to determine the optimal agent configuration. As the next step, the agent cost related parameters β and γ were computed using the prespecified weights w_1 and w_{2i} . Then, the best agent team composition, which maximizes the overall objective $H(s, t)$, is identified from simply searching through the previously generated results. A comparison of the obtained results in terms of the coverage $H(s)$ and the overall objective $H(s, t)$ when the weight w_1 is varied is shown in Fig. 9. The average value of the execution times observed in each method is given in Tab. 8.

Table 8
Observed average execution times.

Method	PGA	BF ₁	BF ₂
Average execution time / (s)	4.56	4328.13	8845.83

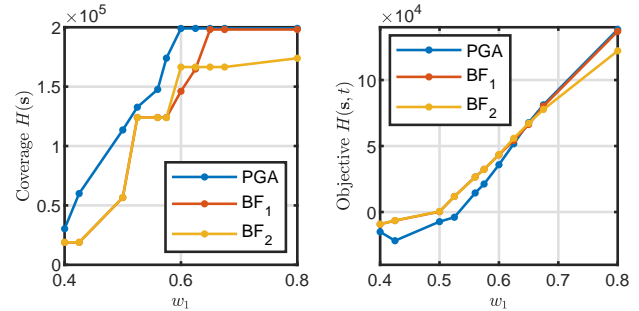
Our main conclusions from this comparison are: (i) The PGA method delivers better coverage levels $H(s)$ across all w_1 values used, and, (ii) As w_1 increases, the PGA method performs better than brute force methods in terms of $H(s, t)$, and, most importantly, (iii) The average execution time required for the PGA method is extremely low compared to brute force approaches (by a factor of 10^{-3}). Finally, we

also emphasize the scalability that the PGA method offers due to its distributed nature.



(a) PGA ($w_1 = 1$): $H(s) = 199037$
Ex.T. = 2.699s
(b) BF₁ (1 Iter.): $H(s) = 198083$
Ex.T. = 246.523s
(c) BF₂ (1 Iter.): $H(s) = 164345$
Ex.T. = 728.374

Fig. 8. Optimal agent configurations, coverage levels, and execution times obtained for the multi-agent coverage problem (see (3)) with 10 heterogeneous agents (see Tab. 4) in a blank mission space using (a) PGA algorithm, (b) Brute force method 1 (BF₁), and, (c) Brute force method 2 (BF₂).



(a) Coverage performance: $H(s)$ (b) Overall performance: $H(s, t)$

Fig. 9. Comparison of coverage performance $H(s)$ and overall performance $H(s, t)$ for different normalization weights w_1 in (11).

6 Conclusions

Multi-agent coverage problem is well-studied when a fixed number of homogeneous agents is to be deployed. In contrast, we address the multi-agent coverage problem where the number of agents to be used is flexible and the available agents are both heterogeneous and have an associated cost value. We have addressed this optimal agent team composition problem by constructing an objective function combining the overall agent team cost with the coverage level delivered by the agent team. An l_1 regularizer is introduced to transform the agent team composition problem into a resource (sensing capacity) allocation problem with no extra non-convexity present. This problem is then solved using a projected gradient ascent (PGA) algorithm initialized through a greedy algorithm and shown to recover the integer-valued variables that were originally relaxed. Further, based on submodularity theory, we have derived tighter performance bounds showing that the PGA algorithm can often lead to near-global-optimal solutions. The effectiveness of the PGA algorithm in diverse mission spaces and heterogeneous multi-agent scenarios has been validated. Additionally, a comparison study with results obtained from a commercial MINLP solver show the efficiency of our proposed

PGA method. An interesting future research direction would be to investigate the applicability of the proposed approach to other multi-agent problems with heterogeneous agents.

The key advantages of the proposed PGA algorithm are: (i) it is capable of solving the combinatorial problem of determining the optimal agent team composition, (ii) it addresses a number of challenges raised due to agent heterogeneity, and (iii) it is characterized by tighter performance bounds for the obtained final solution. Finally we point out that even though the proposed PGA algorithm has been formulated considering a two-dimensional mission space, it is applicable to any N -dimensional mission space upon appropriately modifying the representations of the agent sensing models and the obstacles.

References

- [1] A. Balmann. Modeling land use with multi-agent systems: Perspectives for the analysis of agricultural policies. In *Intl. Institute of Fisheries Economics and Trade Proceedings*, page 11, 2000.
- [2] D. P. Bertsekas. *Nonlinear Programming*. Athena Scientific, 2016.
- [3] A. Breitenmoser, M. Schwager, J.-C. Metzger, R. Siegwart, and D. Rus. Voronoi coverage of non-convex environments with a group of networked robots. In *2010 IEEE Intl. Conf. on Robotics and Automation*, pages 4982–4989, 2010.
- [4] C. H. Caicedo-Nunez and M. Zefran. A coverage algorithm for a class of non-convex regions. In *47th IEEE Conf. on Decision and Control*, pages 4244–4249, 2008.
- [5] C. H. Caicedo-Nunez and M. Zefran. Performing coverage on nonconvex domains. In *IEEE Intl. Conf. on Control Applications*, pages 1019–1024, 2008.
- [6] C. G. Cassandras and W. Li. Sensor networks and cooperative control. In *44th IEEE Conf. on Decision and Control*, volume 2005, pages 4237–4238, dec 2005.
- [7] F. Castanedo, J. García, M. A. Patricio, and J. M. Molina. Data fusion to improve trajectory tracking in a cooperative surveillance multi-agent architecture. *Information Fusion*, 11(3):243–255, 2010.
- [8] M. Conforti and G. Cornuéjols. Submodular set functions, matroids and the greedy algorithm: Tight worst-case bounds and some generalizations of the Rado-Edmonds theorem. *Discrete Applied Mathematics*, 7(3):251–274, 1984.
- [9] J. Cortes, S. Martinez, T. Karatas, and F. Bullo. Coverage control for mobile sensing networks. *IEEE Trans. on Robotics and Automation*, 20(2):243–255, 2004.
- [10] L. Davis. *Handbook of Genetic Algorithms*. London International Thomson Computer Press, Boston, 1996.
- [11] M. L. Fisher, G. L. Nemhauser, and L. A. Wolsey. An analysis of approximations for maximizing submodular set functions—II. *Polyhedral Combinatorics: Dedicated to the memory of D.R. Fulkerson*, pages 73–87, 1978.
- [12] A. Gusrialdi, S. Hirche, T. Hatanaka, and M. Fujita. Voronoi based coverage control with anisotropic sensors. In *American Control Conference*, pages 736–741, 2008.
- [13] A. Gusrialdi and L. Zeng. Distributed deployment algorithms for robotic visual sensor networks in non-convex environment. In *2011 Intl. Conf. on Networking, Sensing and Control*, pages 445–450, 2011.
- [14] Y. Liu, E. K. P. Chong, and A. Pezeshki. Improved bounds for the greedy strategy in optimization problems with curvature. *Journal of Combinatorial Optimization*, 37(4):1126–1149, 2018.
- [15] C. Luo, A. P. Espinosa, D. Pranantha, and A. De Gloria. Multi-robot search and rescue team. In *IEEE Intl. Symp. on Safety, Security, and Rescue Robotics*, pages 296–301, 2011.
- [16] S. Meguerdichian, F. Koushanfar, M. Potkonjak, and M. B. Srivastava. Coverage problems in wireless ad-hoc sensor networks. In *IEEE INFOCOM*, pages 1380–1387, 2001.
- [17] G. L. Nemhauser, L. A. Wolsey, and M. Fisher. An analysis of approximations for maximizing submodular set functions—I. *Mathematical Programming*, 14(1):265–294, 1978.
- [18] M. Schwager, F. Bullo, D. Skelly, and D. Rus. A ladybug exploration strategy for distributed adaptive coverage control. In *2008 IEEE Intl. Conf. on Robotics and Automation*, pages 2346–2353, may 2008.
- [19] X. Sun, C. G. Cassandras, and K. Gokbayrak. Escaping local optima in a class of multi-agent distributed optimization problems: A boosting function approach. In *53rd IEEE Conf. on Decision and Control*, pages 3701–3706, dec 2014.
- [20] X. Sun, C. G. Cassandras, and X. Meng. Exploiting submodularity to quantify near-optimality in multi-agent coverage problems. *Automatica*, 100:349–359, feb 2019.
- [21] R. Tibshirani. Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Society: Series B (Methodological)*, 58(1):267–288, 1996.
- [22] D. Vallejo, J. Albusac, J. J. Castro-Schez, C. Glez-Morcillo, and L. Jiménez. A multi-agent architecture for supporting distributed normality-based intelligent surveillance. *Engineering Applications of Artificial Intelligence*, 24(2):325–340, 2011.
- [23] Z. Wang, B. Moran, X. Wang, and Q. Pan. Approximation for maximizing monotone non-decreasing set functions with a greedy method. *Journal of Combinatorial Optimization*, 31(1):29–43, 2016.
- [24] S. Welikala and C. G. Cassandras. Distributed non-convex optimization of multi-agent systems using boosting functions to escape local optima: Theory and Applications, 2019. [Online]. Available: <http://arxiv.org/abs/1903.04133>.
- [25] S. Welikala and C. G. Cassandras. Distributed Non-convex Optimization of Multi-agent Systems Using Boosting Functions to Escape Local Optima. In *(Accepted) American Control Conference*, 2020.
- [26] Y. Zheng, J. Ma, and L. Wang. Consensus of hybrid multi-agent systems. *IEEE Trans. on Neural Networks and Learning Systems*, 29(4):1359–1365, apr 2018.
- [27] Y. Zheng, Y. Zhu, and L. Wang. Consensus of heterogeneous multi-agent systems. *IET Control Theory and Applications*, 5(16):1881–1888, nov 2011.
- [28] M. Zhong and C. G. Cassandras. Distributed coverage control and data collection with mobile sensor networks. *IEEE Trans. on Automatic Control*, 56(10):2445–2455, 2011.