

# IMAS<sup>2</sup>: Joint Agent Selection and Information-Theoretic Coordinated Perception In Dec-POMDPs

Chongyang Shi  
University of Florida  
Gainesville, United States  
c.shi@ufl.edu

Michael Dorothy  
Army Research Laboratory  
Adelphi, United States  
michael.r.dorothy.civ@army.mil

Wesley A. Suttle  
Army Research Laboratory  
Adelphi, United States  
wesley.a.suttle.ctr@army.mil

Jie Fu  
University of Florida  
Gainesville, United States  
fujie@ufl.edu

## ABSTRACT

We study the problem of jointly selecting sensing agents and synthesizing decentralized active perception policies for the chosen subset of agents within a Decentralized Partially Observable Markov Decision Process (Dec-POMDP) framework. Our approach employs a two-layer optimization structure. In the inner layer, we introduce information-theoretic metrics, defined by the mutual information between the unknown trajectories or some hidden property in the environment and the collective partial observations in the multi-agent system, as a unified objective for active perception problems. We employ various optimization methods to obtain optimal sensor policies that maximize mutual information for distinct active perception tasks. In the outer layer, we prove that under certain conditions, the information-theoretic objectives are monotone and submodular with respect to the subset of observations collected from multiple agents. We then exploit this property to design an IMAS<sup>2</sup> (Information-theoretic Multi-Agent Selection and Sensing) algorithm for joint sensing agent selection and sensing policy synthesis. However, since the policy search space is infinite, we adapt the classical Nemhauser-Wolsey argument to prove that the proposed IMAS<sup>2</sup> algorithm can provide a tight  $(1 - 1/e)$ -guarantee on the performance. Finally, we demonstrate the effectiveness of our approach in a multi-agent cooperative perception in a grid-world environment.

## KEYWORDS

Decentralized Partially Observable Markov Decision Process (Dec-POMDP); Agent Selection; Active Perception; Submodular Optimization.

### ACM Reference Format:

Chongyang Shi, Wesley A. Suttle, Michael Dorothy, and Jie Fu. 2026. IMAS<sup>2</sup>: Joint Agent Selection and Information-Theoretic Coordinated Perception In Dec-POMDPs. In *Proc. of the 25th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2026)*, Paphos, Cyprus, May 25 – 29, 2026, IFAAMAS, 9 pages. <https://doi.org/10.65109/STZK9664>



This work is licensed under a Creative Commons Attribution International 4.0 License.

*Proc. of the 25th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2026)*, C. Amato, L. Dennis, V. Mascardi, J. Thangarajah (eds.), May 25 – 29, 2026, Paphos, Cyprus. © 2026 International Foundation for Autonomous Agents and Multiagent Systems ([www.ifaamas.org](http://www.ifaamas.org)). <https://doi.org/10.65109/STZK9664>

## 1 INTRODUCTION

Autonomous multi-agent systems are increasingly deployed in environments where agents must actively gather information under uncertainty. Applications include teams of robots for surveillance and search-and-rescue [7, 17, 18], sensor networks for target tracking [6, 13], and cooperative perception in autonomous driving and aerial monitoring [24, 25]. In such settings, an active perception agent must decide not only how to act, but also what to sense, given the past observation. In multi-agent cooperative active perception, it must additionally account for its knowledge of other agents' possible observations to avoid redundancy and achieve better coordination. Meanwhile, in addition to perception, a multi-agent system may be required to perform tasks leveraging the information from the perception agents, for example, in a multi-UAV target tracking problem, a subset of UAVs may be tasked to ensure accurate tracking of the moving target, while other UAVs may be tasked to interdict the target given the target trajectory.

Motivated by these applications, this paper studies the following research problem: "Given a multi-agent system with heterogeneous dynamics and perception capabilities, how to select a subset of agents and design their decentralized perception strategies for a perception objective?" Specifically, we model a multi-agent system and its interaction with the dynamic, stochastic environment using a decentralized partially observable Markov decision process (Dec-POMDP) framework. We aim to select a subset of the agents for perception tasks, whose objectives are defined by maximizing the mutual information between some unknown quantity (trajectory, states, or critical events) and the collective observations of the perception team.

**Related work.** This problem is closely related to the sensor placement problem, which aims to find an optimal subset of agents to participate in a perception task under limited sensing resources. Unlike traditional sensor coverage problems that emphasize spatial reach or visibility, the goal is to maximize the informativeness of the selected agents' observations about the environment or latent variables of interest. Selecting too many agents wastes resources and increases redundancy, while selecting too few may lead to poor estimation or inference accuracy. To balance these trade-offs, many existing approaches formulate the objective using information-theoretic criteria, such as mutual information or entropy reduction [9, 23]. When the agents' policies are finite, the environment is

deterministic, and the inference objective is stationary—as in multi-robot path planning for predicting algae content in a lake—the resulting objectives often exhibit submodularity. This property enables the design of greedy selection algorithms with provable near-optimal performance guarantees and strong scalability for large multi-agent systems. However, in stochastic multi-agent systems, where the policy space is infinite (often parameterized by deep neural networks) and the unknown quantities to be inferred can be a stochastic process (such as tracking a moving agent), existing results do not directly apply.

The problem of decentralized active sensing and perception has been studied extensively in multi-agent systems. Early work, such as [10], formulated sensor management as a centralized active sensing problem, where a single decision maker plans sensing actions to maximize expected information gain using Bayesian filtering. While effective for small-scale systems, such centralized formulations are computationally expensive and scale poorly with the number of agents. Subsequent research has focused on decentralized and scalable planning methods. Satsangi et al. [19, 20] studied a class of partially observable Markov decision processes (POMDPs) whose value functions exhibit submodularity, and leveraged this property to design efficient algorithms for dynamic sensor selection with near-optimal guarantees. In contrast, Kumar et al. [11] considered multi-agent planning under stochastic dynamics, where submodularity appears in the reward structure rather than the value function. Their framework provides a theoretical foundation for decentralized decision making with submodular rewards, enabling agents to coordinate without centralized control. Best et al. [2] proposed Dec-MCTS, a decentralized Monte Carlo tree search method, for multi-robot active perception, which can handle stochasticity, but does not exploit the submodular structure of reward functions for efficiency or near-optimality guarantees. Lauri and Oliehoek [12] introduced a prediction-reward framework that quantifies the uncertainty in a centralized state estimate obtained after all agents complete sensing. To enable decentralized computation, they approximate this centralized reward using the expected accuracy of each agent’s prediction of the final belief. This reformulation expresses the global information objective as a standard Dec-POMDP reward function dependent only on local states, actions, and prediction actions, allowing the use of existing Dec-POMDP solvers while preserving the informativeness objective. However, this method also does not employ submodularity for more efficient computation.

Recent research also studied the multi-robot multi-target tracking problem in deterministic, nonlinear systems [3]. They employ the receding horizon planning framework and consider at each planning epoch, the team of robots aims to optimize the mutual information between the target states and the observations generated by a set of feasible trajectories for each robot. Because they consider a finite set of trajectories for each robot and deterministic dynamics, such an objective function can be shown to be submodular and monotone. By leveraging the property of a submodular function, they develop an algorithm to efficiently compute the near-optimal trajectories for multiple robots.

**Our contributions.** Existing work on sensor selection or motion planning of active sensing agents considers finitely many sensors [23] or discrete trajectories [3]. In our setting, though the set of

agents is finite, each agent is to compute an observation-based stochastic policy, and therefore, the policy space for agents is infinite. Therefore, existing GreedyMax algorithms for submodular [16] cannot be directly applied to compute jointly the subset of agents and their decentralized perception policies. Our problem can not be directly mapped to decentralized planning for multi-agent systems given submodular rewards because the information objectives cannot be decomposed into submodular reward given the joint state and joint actions (see Section 3). Existing Dec-POMDP for active perception [2, 12] does not solve the agent selection problem jointly with the policy synthesis.

First, we prove that in a Dec-POMDP, when the set of policies are fixed, then the mutual information between the joint state trajectory given a subset of agents’ policies, and consequently their partial observations, is monotone and submodular, provided that the agents’ observations are conditionally independent given the joint trajectory. Building on this result, we then investigate two other inference objectives: (1) Inferring the trajectory of an environment agent; (2) and inferring a secret property which is a function of the trajectory of an environment agent. We show that in both cases, under certain assumption, the mutual information between the quantity to be inferred and a subset of agents’ observations are submodular and monotone.

However, these results only enable us to use approximate algorithms for submodular maximization to solve the agent selection problem, assuming each selected agent follows a pre-defined decentralized policy. To achieve joint agent selection and policy synthesis, we propose the IMAS<sup>2</sup> (Information-theoretic Multi-Agent Selection and Sensing) algorithm, a variant of the GreedyMax algorithm [16] to select one agent and determine its perception policy at a time based on the principle of maximizing the marginal gain. We then prove that, under additional constraints on subsequent maximal marginal gains, this algorithm can provide a strong  $(1 - 1/e)$ -guarantee on the performance. The proposed algorithm achieves scalability and near-optimal performance, bridging rigorous submodular optimization theory with practical decentralized planning in multi-agent active perception. Finally, we formulate the agent selection and policy optimization problem as a two-layer optimization: an inner layer that computes an optimal perception policy, with respect to maximizing information gain of a chosen agent, in addition to observations of selected agents. And an outer layer efficiently selects  $k$  agents to achieve the approximately optimal cooperative active perception. The results are then experimentally validated.

**Assumptions and Scope.** We focus on Dec-POMDPs with observation independent sensing models, formalized by Assumption 1. Importantly, Subsection 3.1 (Inferring Latent State Sequence) only requires this observation conditional independence and does not require transition independence. In contrast, Subsection 3.2 (Inferring Environment State Sequence) and Section 3.3 (Environment Secret Estimation) adopt the stronger transition- and observation independence conditions in Assumption 2. We highlight these assumptions explicitly to distinguish the settings covered by each result.

## 2 PRELIMINARY

**Notations.** The set of real numbers is denoted by  $\mathbb{R}$ . Random variables will be denoted by capital letters, and their realizations by lowercase letters (e.g.,  $X$  and  $x$ ). A sequence of random variables and their realizations with length  $T$  are denoted as  $X_{0:T}$  and  $x_{0:T}$ . The notation  $x_i$  refers to the  $i$ -th component of a vector  $x \in \mathbb{R}^n$  or to the  $i$ -th element of a sequence  $x_0, x_1, \dots$ , which will be clarified by the context. Given a finite set  $\mathcal{S}$ , let  $\mathcal{D}(\mathcal{S})$  be the set of all probability distributions over  $\mathcal{S}$ . The set  $\mathcal{S}^T$  denotes the set of sequences with length  $T$  composed of elements from  $\mathcal{S}$ , and  $\mathcal{S}^*$  denotes the set of all finite sequences generated from  $\mathcal{S}$ . The empty string in  $\mathcal{S}^*$  is denoted by  $\emptyset$ . The notation  $X_1 \perp\!\!\!\perp X_2|Y$  means random variables  $X_1$  and  $X_2$  are conditionally independent given random variable  $Y$ .

We model the interaction between a multi-agent system and its environment using a finite-horizon, decentralized, partially observable Markov decision process (Dec-POMDP) without reward:

$$\mathcal{M} = \langle T, \mathcal{N}, \mathcal{S}, \mathcal{A}, \mathcal{O}, P, \mu, \{E_i\} \rangle,$$

where:

- $T \in \mathbb{N}$  is the time horizon of the problem;
- $\mathcal{N} = \{1, 2, \dots, N\}$  is a set of  $N$  agents,
- $\mathcal{S}$  is the finite set of states.
- $\mathcal{A} = \prod_{i \in \mathcal{N}} \mathcal{A}_i$  is the joint action space, where  $\mathcal{A}_i$  is the finite action space of agent  $i \in \mathcal{N}$ . The tuple  $a = \langle a_1, a_2, \dots, a_N \rangle$  is called the joint action.
- $\mathcal{O} = \prod_{i \in \mathcal{N}} \mathcal{O}_i$  is the joint observation space, where  $\mathcal{O}_i$  is the finite observation space of agent  $i$ . The tuple  $o = \langle o_1, o_2, \dots, o_N \rangle$  of individual observations is called the joint observation;
- $P : \mathcal{S} \times \mathcal{A} \rightarrow \mathcal{D}(\mathcal{S})$  is the probabilistic transition function;
- $\mu$  is the initial state distribution.
- $E_i : \mathcal{S} \rightarrow \mathcal{D}(\mathcal{O}_i)$  is the emission/observation function for agent  $i$ .

An admissible solution of a Dec-POMDP is a decentralized joint policy  $\pi$ , i.e., a tuple  $\langle \pi_1, \dots, \pi_N \rangle$ , where the individual policy  $\pi_i : (\mathcal{O}_i \times \mathcal{A}_i)^* \mathcal{O}_i \rightarrow \mathcal{D}(\mathcal{A}_i)$  of each agent  $i$  maps individual **observation-action** sequences  $y_{i,0:t} = (o_{i,0}, a_{i,0}, o_{i,1}, \dots, o_{i,t})$  to an individual action distribution. It is noted that the individual agent knows its own action but not others' actions.

We introduce the notion of submodular functions next. Denote  $\Omega$  as a ground set of  $n$  data points,  $\Omega = \{x_1, x_2, x_3, \dots, x_n\}$ , and a set function  $f : 2^\Omega \rightarrow \mathbb{R}$ . We say that  $f$  is *normalized* if  $f(\emptyset) = 0$ , and  $f$  is *subadditive* if

$$f(U) + f(V) \geq f(U \cup V),$$

holds for all  $U, V \subset \Omega$ . Define the first-order partial derivative (equivalently, the gain) of an element  $j \notin U$  in the context  $U$  as

$$f(j | U) = f(U \cup \{j\}) - f(U).$$

The function  $f$  is *submodular* [5] if for all  $U, V \subset \Omega$ , it holds that

$$f(U) + f(V) \geq f(U \cup V) + f(U \cap V).$$

An equivalent characterization of submodularity is the *diminishing marginal returns* property, namely

$$f(j | U) \geq f(j | V) \quad \text{for all } U \subset V, j \notin V.$$

Next, we introduce the definition of entropy and mutual information. Entropy is commonly employed to quantify the uncertainty

about a random variable [8]. The conditional entropy of a random variable  $X_2$  given another random variable  $X_1$  is defined by

$$H(X_2|X_1) = - \sum_{x_1 \in \mathcal{X}} \sum_{x_2 \in \mathcal{X}} p(x_1, x_2) \log p(x_2|x_1).$$

Conditional entropy quantifies the uncertainty of  $X_2$  given  $X_1$ , and lower values indicate that  $X_2$  is easier to infer from knowing the value of  $X_1$ . The mutual information between two random variables  $X_1$  and  $X_2$  is defined by

$$I(X_1; X_2) = \sum_{x_1 \in \mathcal{X}} \sum_{x_2 \in \mathcal{X}} p(x_1, x_2) \log \frac{p(x_1, x_2)}{p(x_1)p(x_2)}.$$

Mutual information quantifies the amount of information that one random variable contains about another. Higher values indicate a stronger statistical dependency between  $X_1$  and  $X_2$ , while  $I(X_1; X_2) = 0$  if and only if  $X_1$  and  $X_2$  are independent.

## 3 LEVERAGING SUBMODULARITY FOR MULTI-AGENT ACTIVE PERCEPTION

This section examines which classes of inference objectives in a multi-agent active perception planning exhibit submodularity with respect to the set of agents' observations.

### 3.1 Inferring Latent State Sequence

In this subsection, we study a class of active perception objectives for minimizing the uncertainty in the estimation of the joint state trajectories. In particular, we are interested in finding an (approximately)-optimal subset  $\mathcal{K}$  of the agents to carry out the active perception tasks and computing their individual policies, for some  $\mathcal{K} \subset \mathcal{N}$ . The agents in  $\mathcal{N} \setminus \mathcal{K}$  are excluded from the perception task and their observations do not contribute to the perception objective. Importantly, selecting a subset  $\mathcal{K}$  does *not* remove the other agents from the system dynamics: all agents in  $\mathcal{N}$  still execute their respective policies and evolve under the original joint transition model. The only change is that the inference objective is evaluated using the collective observations of agents in  $\mathcal{K}$  (e.g., for  $\mathcal{N} = \{1, \dots, 5\}$  and  $\mathcal{K} = \{1, 2, 4\}$ , all five agents act and affect the state evolution, while only  $Y_1, Y_2, Y_4$  are used for estimating  $S_{0:T}$ ).

**PROBLEM 1.** *Given the Dec-POMDP  $\mathcal{M}$ , determine a subset  $\mathcal{K}$  of  $k$  agents and design the joint policy  $\{\pi_i, i \in \mathcal{K}\}$  that maximizes the mutual information between the latent state sequence  $X := S_{0:T}$  and the observations:*

$$\underset{\mathcal{K} \subset \mathcal{N}, |\mathcal{K}|=k, \pi_{\mathcal{K}}=\{\pi_k, k \in \mathcal{K}\}}{\text{maximize}} \quad I(X; Y_{\mathcal{K}}, M_{\pi_{\mathcal{K}}}) \quad (1)$$

where  $\pi_{\mathcal{K}}$  is the joint policy for the selected subset of agents, and  $M_{\pi_{\mathcal{K}}}(\mathcal{K})$  is the induced stochastic process given the joint policy  $\pi_{\mathcal{K}}$  for the selected agents in  $\mathcal{K}$  and arbitrary policies for the non-selected agents in  $\mathcal{N} \setminus \mathcal{K}$ ;  $Y_{\mathcal{K}}$  is the collective observations of the selected  $k$  agents.

The following assumption is made to establish the submodularity in the perception objective with respect to agents' observations.

**ASSUMPTION 1.** *For any  $s \in \mathcal{S}$ , for any  $i, j \in \mathcal{N}$  where  $i \neq j$ , agent  $i$ 's observation  $E_i(s)$  is conditionally independent from agent  $j$ 's observation  $E_j(s)$  given state  $s$ .*

In the following, we fix a joint policy  $\pi$ , consisting of individual agents' observation-based policies. Let the induced stochastic process be  $M_\pi = \{S_t, \{A_{i,t}, i \in \mathcal{N}\}, \{O_{i,t}, i \in \mathcal{N}\}, t \in \mathcal{N}\}$ . We denote by  $Y_i := O_{i,0:T}$  the observation sequence received by agent  $i$ . The choice of this joint policy  $\pi$  is arbitrary, i.e., it need not be the optimal policy for a given perception task.

LEMMA 1. *Given any  $Y_i = O_{i,0:T}, Y_j = O_{j,0:T}$  representing agents  $i$  and  $j$ 's observation sequences for a finite horizon  $T$  and  $X = S_{0:T}$  be the latent state sequence. Under Assumption 1,  $Y_i$  and  $Y_j$  are conditional independent given  $X$ .*

PROOF. See appendix<sup>1</sup>. □

LEMMA 2. *Let  $A \subset \mathcal{N}$  and  $Y_A = \{Y_i : i \in A\}$ . For any  $Y_j$  where  $j \notin A, H(Y_j|Y_A, X) = H(Y_j|X)$ .*

PROOF. Due to the conditional independence between  $Y_j$  and any  $Y_i \in Y_A$  given  $X$ . □

LEMMA 3. *Let  $A \subset \mathcal{N}$  and  $Y_A = \{Y_i : i \in A\}$ . Let*

$$g(Y_A) := I(X; Y_A),$$

where  $I(\cdot; \cdot)$  denotes mutual information between the random state sequence  $X$  and the collection of observations  $Y_A$ . The function  $g(\cdot)$  is monotone and submodular.

PROOF. See appendix. □

Under the assumption of independent observations but possibly coupled dynamics, we have proven that the mutual information between the latent state sequence  $X$  and a subset  $Y_A$  of observations is monotone submodular in the subset  $Y_A$ .

REMARK 1. *The above analysis is based on a pre-defined decentralized policy. It can be used for selecting a subset of agents whose observations are most informative for estimating the joint state trajectory. Leveraging the monotone, submodular property, the original GreedyMax algorithm [16] can find such a subset of agents with  $(1 - 1/e)$  performance guarantee on the suboptimality.*

### 3.2 Inferring Environment State Sequence

Although mutual information is often used to quantify information gain, directly maximizing the mutual information between the joint state trajectory and a subset of perception agents' observations may not lead to the intended inference objective. Since  $I(X; Y) = H(X) - H(X|Y)$ , maximizing mutual information encourages reducing the uncertainty of the trajectory given observations, measured by  $H(X|Y)$ . However, it also rewards increasing the prior entropy  $H(X)$ . As a result, this objective can favor policies that induce joint state trajectories with a high initial uncertainty, rather than those that genuinely improve inference accuracy.

Therefore, we consider another objective function that is well-suited for environmental monitoring tasks. We focus on a case where a team of agents is deployed to actively monitor a dynamic environment state trajectory. In the Dec-POMDP, each joint state  $s = \langle s_e, s_1, \dots, s_N \rangle$  consists of an environment agent's state  $s_e$  and agent  $i$ 's state  $s_i$  for each  $i \in \mathcal{N}$  of the multi-agent system.

We modify the objective function in (1) for the environment-state trajectory estimation as follows:

$$\underset{\mathcal{K} \subset \mathcal{N}, |\mathcal{K}|=k, \pi_{\mathcal{K}} = \{\pi_k, k \in \mathcal{K}\}}{\text{maximize}} I(X_e; Y_{\mathcal{K}}, M_{\pi_{\mathcal{K}}}) \quad (2)$$

where  $X_e = S_{e,0:T}$  is the environment state trajectory.

Since the environment state sequence is uncontrollable,  $H(X_e)$  is a constant. In this case, given  $I(X; Y) = H(X) - H(X|Y)$ , maximizing the mutual information is equivalent to minimizing the uncertainty in the environment state trajectory. That is, the optimization problem in (2) is equivalent to

$$\underset{\mathcal{K} \subset \mathcal{N}, |\mathcal{K}|=k, \pi_{\mathcal{K}} = \{\pi_k, k \in \mathcal{K}\}}{\text{maximize}} -H(X_e|Y_{\mathcal{K}}, M_{\pi_{\mathcal{K}}}) \quad (3)$$

In general, the objective in (2) is not submodular in the set of agents' observations. This is because of the local observations of two agents may no longer be conditionally independent given the environment-state trajectory. However, we show that for a subset of the problems, we can also derive submodularity in (2).

Moreover, under the transition- and observation-independence conditions below, the Dec-POMDP induced by selecting a subset of agents can be constructed in a standard way by restricting the dynamics and observation model to the environment state and the selected agents' local states. In particular, because the joint transition and observation models factor across agents, the evolution of the selected agents depends only on their own local states/actions and the environment process, and the non-selected agents can be marginalized without affecting the resulting reduced model.

ASSUMPTION 2. *The agents and their environment have independent dynamics. That is, there exists transition functions  $P_e, \{P_i, i \in \mathcal{N}\}$  such that*

$$\begin{aligned} P((s'_e, s'_1, \dots, s'_N)|(s_e, s_1, \dots, s_N), (a_1, \dots, a_N)) \\ = P_e(s'_e|s_e) \prod_{i=1}^N P_i(s'_i|s_i, a_i). \end{aligned} \quad (4)$$

The initial state distribution can also be decomposed as

$$\mu_0((s_e, s_1, \dots, s_N)) = \mu_{0,e}(s_e) \cdot \prod_{i=1}^N \mu_{0,i}(s_i).$$

The observation of agent  $i$  is independent from the other agents' states: That is, if  $s_e = s'_e$  and  $s_i = s'_i$ , then for any  $s_j, s'_j$  when  $j \neq i$ ,

$$E_i(s_e, s_1, \dots, s_N) = E_i(s'_e, s'_1, \dots, s'_N).$$

In this case, we define a local observation function  $\hat{E}_i : S_e \times S_i \rightarrow \mathcal{D}(O_i)$  such that  $\hat{E}_i(s_e, s_i) := E_i(s_e, s_1, \dots, s_N)$ .

Due to the decoupled dynamics and observation, we can define individual agent's observation sequence based on its own policy. Let  $\pi_i$  be an observation-dependent policy for agent  $i$ . The  $\pi_i$ -induced local observation  $Y_i = O_{i,0:T}$  is defined by

$$\Pr(O_{i,t} = o) = \hat{E}_i(o|S_{e,t}, S_{i,t}),$$

where  $S_{e,t} \sim P_e(\cdot|S_{e,t-1}), S_{i,t} \sim P_i(\cdot|S_{i,t-1}, A_{i,t}), S_{i,0} \sim \mu_{0,i}(\cdot)$  and  $A_{i,t} \sim \pi_i(\cdot|O_{i,0:t}, A_{i,0:t-1})$ , for any  $t \in \{1, \dots, T\}$ .

The following result again assumes a fixed joint policy  $\pi = \langle \pi_1, \dots, \pi_N \rangle$  and the induced stochastic process of states, actions, and observations.

<sup>1</sup>You can find the appendix in <https://github.com/AronYoung414/multi-agent-active-perception-grid-world>.

LEMMA 4. For any  $i, j \in \mathcal{N}, i \neq j$ . Let  $Y_i, Y_j$  be the local observation sequences for agents  $i, j$ . Under Assumption 2, let  $X_e := S_{e,0:T}$  be the latent state sequence of the environment agent, then  $Y_i$  and  $Y_j$  are conditionally independent given  $X_e$ .

PROOF. See appendix.  $\square$

COROLLARY 1. Let  $A \subset \mathcal{N}$  and  $Y_A = \{Y_i : i \in A\}$ . For any  $Y_j$  where  $j \notin A$ ,  $H(Y_j|Y_A, X_e) = H(Y_j|X_e)$ .

PROOF. The proof is similar to that of Lemma 2 given Lemma 4.  $\square$

LEMMA 5. Let  $A \subset \mathcal{N}$  and  $Y_A = \{Y_i : i \in A\}$ . Let  $g_e(Y_A) := I(X_e; Y_A)$ , where  $I(X_e; Y_A)$  is the mutual information between the environmental state sequence  $X_e$  and the collective observations  $Y_A$ . The function  $g_e(\cdot)$  is monotone and submodular.

PROOF. Similar to the proof of Lemma 3 where we use Lemma 4.  $\square$

### 3.3 Environment Secret Estimation

In this section, we consider a perception objective besides the aforementioned environment trajectory estimation. As in Subsection 3.2, our analysis for this setting assumes the transition- and observation-independence conditions in Assumption 2, so that the Dec-POMDP induced by a reduced agent set can be constructed by restricting to the environment state and the selected agents' local states/observations. Consider a random variable  $Z$  defined as a surjective function of  $X_e$  i.e.,  $Z = f(X_e)$ . It is noted that  $H(X_e)$  is a constant and hence  $H(Z)$  is a constant.

The agent selection and optimal perception planning problem for inferring the value of  $Z$  given the collective observations is formulated as the following optimization problem:

$$\max_{\mathcal{K} \subset \mathcal{N}, |\mathcal{K}|=k, \pi_{\mathcal{K}} = \{\pi_k, k \in \mathcal{K}\}} -H(Z|Y_{\mathcal{K}}, M_{\pi_{\mathcal{K}}}) \quad (5)$$

First, based on the relation between mutual information and entropy, we have  $I(Z; Y_A) = H(Z) - H(Z|Y_A)$ , and thus

$$I(X_e; Y_A) - I(Z; Y_A) = H(X_e) - H(X_e|Y_A) - H(Z) + H(Z|Y_A).$$

Given the fact that  $H(X_e|Y_A) - H(Z|Y_A) = H(X_e|Z, Y_A)$ , we have

$$I(Z; Y_A) = I(X_e; Y_A) - H(X_e) + H(Z) + H(X_e|Z, Y_A). \quad (6)$$

We will show that  $I(Z; Y_A)$  is monotone and submodular in the following lemma.

LEMMA 6 (APPROXIMATE SUBMODULARITY). Let

$$g(A) := I(Z; Y_A),$$

where  $Z$  is a secret variable and  $Y_A$  denotes the joint observations collected by agent set  $A$ . Then  $g(\cdot)$  is monotone and  $\epsilon$ -approximately submodular, i.e., there exists a submodular function

$$h(A) := I(X_e; Y_A)$$

such that for all  $A \subseteq \mathcal{N}$ ,

$$(1 - \epsilon)h(A) \leq g(A) \leq (1 + \epsilon)h(A),$$

where

$$\epsilon := \max_A \frac{I(X_e; Y_A | Z)}{I(X_e; Y_A)} \in (0, 1].$$

---

### Algorithm 1 IMAS<sup>2</sup> Algorithm

---

**Require:** Ground set of agents  $\mathcal{N}$ , budget  $k$

**Ensure:** Selected agent set  $\mathcal{K}$

- 1: Initialize  $\mathcal{K}^{(0)} \leftarrow \emptyset, \pi^{(0)} \leftarrow \emptyset$ .
  - 2: **for**  $i = 1$  to  $k$  **do**
  - 3:   Select  $(j^*, \pi_{j^*}) = \operatorname{argmax}_{j \in \mathcal{N} \setminus \mathcal{K}^{(i-1)}} \max_{\pi_j \in \Pi_j} \left( I(X; Y_{\mathcal{K}^{(i-1)} \cup \{j\}}, M_{\pi_{\mathcal{K}^{(i-1)} \cup \{j\}}}) - I(X; Y_{\mathcal{K}^{(i-1)}}, M_{\pi_{\mathcal{K}^{(i-1)}}}) \right)$
  - 4:   Update  $\mathcal{K}^{(i)} \leftarrow \mathcal{K}^{(i-1)} \cup \{j^*\}, \pi^{(i)} \leftarrow \pi^{(i-1)} \cup \{\pi_{j^*}\}$
  - 5: **end for**
  - 6: **return**  $\mathcal{K} := \mathcal{K}^{(k)}, \pi^* := \pi^{(k)}$ .
- 

PROOF. See appendix.  $\square$

## 4 APPROXIMATION ALGORITHMS FOR JOINT AGENT SELECTION AND POLICY SYNTHESIS

The previous section shows the submodularity for several inference objectives under various assumptions for Dec-POMDP, while the joint policies are assumed to be fixed. Thus, these results do not solve the problems in (1), (2), or (5), since they all require the joint agent selection and policy synthesis.

### 4.1 Approximation Algorithm and Performance Guarantee

In this section, we devise an efficient method called IMAS<sup>2</sup> (Algorithm 1) to solve these problems and analyze the performance guarantee. The variable  $X$  to be inferred can be changed to  $X_e$  or  $Z$ , depending on (1), (2), or (5).

In each iteration of Algorithm 1, for each candidate agent, its optimal local policy is computed to maximize the mutual information between the agent's observations and the latent variables of interest. Then, with the computed policies, the algorithm evaluates the marginal gain contributed by each agent and greedily selects the next one that yields the largest improvement.

A key property of submodular functions is that they admit strong approximation guarantees when optimized under cardinality constraints. However, because we are simultaneously selecting the agent and determining their policy, the proof of approximate optimality needs to be modified. Next, we derive the performance guarantee of the IMAS<sup>2</sup> algorithm when the objective functions are submodular.

In (1), let  $f(A, \pi_A) = I(X; Y_A, M_{\pi_A})$  and in (2), let  $f(A, \pi_A) = I(X_e; Y_A, M_{\pi_A})$ , for any subset  $A \subset \mathcal{A}$  and  $\pi_A = \{\pi_i, i \in A\}$ . Let  $(\mathcal{K}^*, \pi^*)$  be the optimal solution to (1) (or (2)). Let  $(\mathcal{K}^{(i)}, \pi^{(i)})$  be the solution returned by Algorithm 1 after the  $i$ -th selection. By default,  $\mathcal{K}^{(0)} = \emptyset$  and  $\pi^{(0)} = \emptyset$ . By default,  $f(\emptyset, \emptyset) = 0$ .

PROPOSITION 1. Let  $\Delta_i := f(\mathcal{K}^{(i)}, \pi^{(i)}) - f(\mathcal{K}^{(i-1)}, \pi^{(i-1)})$  for  $1 \leq i < k$ . It holds that  $\Delta_{i+1} \leq \Delta_i, \forall 0 \leq i < k$ .

PROOF. See appendix.  $\square$

**THEOREM 1.** Assuming for all  $0 \leq i < k$ ,  $\frac{\Delta_i}{\Delta_{i+1}} \leq \frac{k+1}{k}$ , then

$$f(\mathcal{K}^{(k)}, \boldsymbol{\pi}^{(k)}) \geq (1 - \frac{1}{e})f(\mathcal{K}^*, \boldsymbol{\pi}^*). \quad (7)$$

**PROOF.** See appendix. □

**REMARK 2.** Since the objective function in (5) is only approximately submodular, analyzing the performance of the proposed algorithm requires explicitly accounting for the approximation ratio  $\epsilon$ . We defer a performance analysis and the establishment of formal guarantees for the secret inference case to future work.

### 4.2 Policy Synthesis for Maximizing Information Gain

In Algorithm 1, the outer maximization is straightforward: the optimal agent is selected by comparing all candidate agents. However, we must also address the inner policy optimization problem: for each  $j \in \mathcal{N} \setminus \mathcal{K}^{(i)}$ ,

$$\boldsymbol{\pi}_{j^*} = \arg \max_{\boldsymbol{\pi}_j} (I(X; Y_{\mathcal{K} \cup \{j\}}, M_{\boldsymbol{\pi}_{\mathcal{K} \cup \{j\}}}) - I(X; Y_{\mathcal{K}}, M_{\boldsymbol{\pi}_{\mathcal{K}}}))$$

We discuss how to solve the optimization problem using existing results in literature.

*Late state sequence estimation:* In [15], the authors develop a method to compute an optimal policy that maximizes  $I(X; Y, M_\theta)$  for a single-agent POMDP. This method can be used for solving the inner maximization problem: First, we compute the single-agent POMDP for agent  $j$  by fixing the selected agents' policies  $\boldsymbol{\pi}_{\mathcal{K}^{(i)}}$  computed from the previous iteration; second, the algorithm for trajectory estimation in single-agent POMDP [15] can be applied. Note that the observations received by selected agents are included in the observation of the single-agent POMDP.

*Inferring environment state trajectory or secret.* The method in [22] employs a policy gradient approach. For agent  $i$ , let  $\{\pi_{i,\theta} | \theta_i \in \Theta\}$  be a class of parameterized stochastic policies. We denote the joint policy parameter for all agents as  $\boldsymbol{\theta}$  and the corresponding joint policy as  $\boldsymbol{\pi}_\theta$ . Let  $Y(\theta_j)$  be agent  $j$ 's observation under agent  $j$ 's policy parameterized by  $\theta_j$  and  $Y(\boldsymbol{\theta}_{\mathcal{K}^{(i)}})$  is the joint observation of agents in  $\mathcal{K}^{(i)}$  given the joint policy of selected agents.


As argued in the section 3.2, maximizing the mutual information for these two cases is equivalent to minimizing the conditional entropy in the unknown variable. To obtain the locally optimal policy parameter  $\theta_j$ , we initialize a policy parameter  $\theta_j$  and carry out the gradient descent. At each iteration  $\tau \geq 0$ , let  $\boldsymbol{\theta} := \boldsymbol{\theta}_{\mathcal{K}^{(i)}}$  represent the joint policy of selected agents for notation simplicity,

$$\theta_j^{\tau+1} = \theta_j^\tau - \eta [\nabla_{\theta_j} H(X_e | Y(\boldsymbol{\theta}) \cup Y(\theta_j)) |_{\theta_j = \theta_j^\tau}], \quad (8)$$

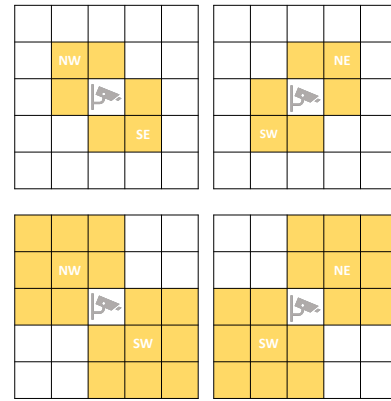
where  $\eta$  is the step size (learning rate). The gradient of conditional entropy is given by

$$\nabla_{\theta_j} H(X_e | Y(\boldsymbol{\theta})) = \mathbb{E}_{\mathbf{y} \sim \mathcal{M}_\theta} [H(X_e | Y = \mathbf{y}) \nabla_{\theta_j} \log P_\theta(y_j)]. \quad (9)$$

Notably, when evaluating the conditional entropy, we use the joint observation  $\mathbf{y}$  including both the observations of selected agents and the agent  $j$ 's observation under the current policy  $\theta_j^\tau$ . When we are evaluating the gradient of the log probability of observations, we only use the observations of agent  $j$  because the other selected



**Figure 1: The 10 × 10 grid world environment.**



**Figure 2: The sensor range of different actions (Top: small range sensors; Bottom: large range sensors).**

agents' policies are fixed and independent from  $\theta_j$ . We have the following result from [22]:

$$\nabla_{\theta_j} \log P_\theta(y_j) = \sum_{t=0}^T \nabla_{\theta_j} \log \pi_{\theta_j}(a_t | o_{j,0:t-1}), \quad (10)$$

where  $o_{j,0:t-1}$  is the local observation of agent  $j$ . The policy gradient from minimizing  $H(Z|Y(\boldsymbol{\theta}))$  is derived analogously (see [22] for technical details.)

## 5 EXPERIMENT

We illustrate the effectiveness of our IMAS<sup>2</sup> algorithm using a multi-agent selection and sensing strategy design problem in a 10 × 10 grid-world environment (Figure 1). The multi-agent system is a directed sensor network. The underlying state of the environment includes a tuple  $s_e = (x, y, \mathbf{b})$  where  $x, y$  are the position of an environment agent, marked by the robot, which are not controlled. The last variable  $\mathbf{b}$  is a Boolean variable. If  $\mathbf{b} = 0$ , then the robot

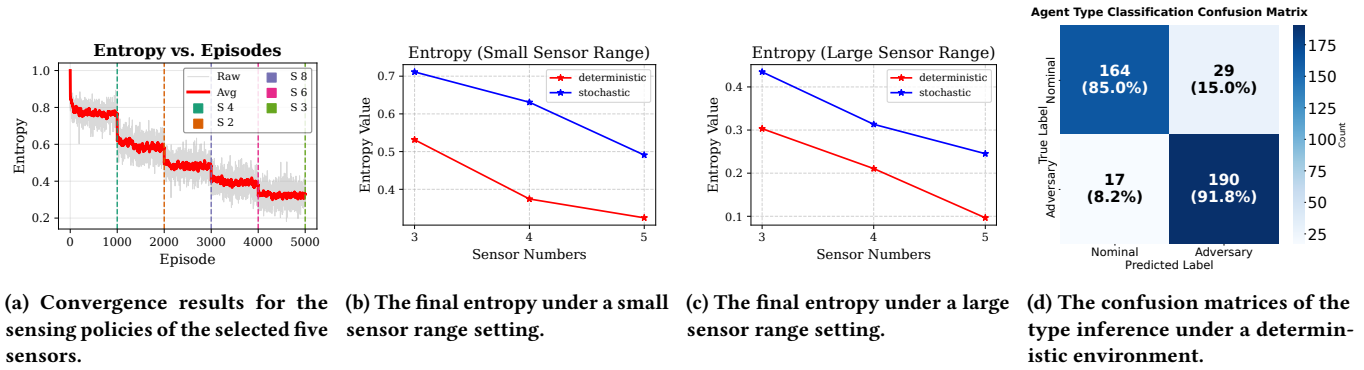


Figure 3: Comparison of convergence, entropy, and inference performance under different sensing conditions.

is benign and follows a stochastic shortest path policy to reach the “Normal goal” (marked by red, filled flag). If  $\mathbf{b} = 1$ , then the robot is adversarial and follows a stochastic shortest path policy to reach the “Adversary’s goal” (marked by red, unfilled flag). Once a robot reaches its goal, it remains there. The boundary and obstacles are bouncing obstacles. The robot can move north, south, east, or west, or choose to stay in place. The initial position of the robot is uniformly distributed on the first column of the grid world, and the type of the robot is uniformly sampled from a Bernoulli distribution for  $\mathbf{b}$  with  $p = 0.5$ .

In the later experiments, we consider two types of robot dynamics: deterministic and stochastic. Under deterministic dynamics, an action directly results in the intended movement (e.g., moving north always results in a northward transition). Under stochastic dynamics, however, when an agent attempts to move in a given direction, there is a 0.1 probability of deviating to one of the orthogonal directions. For example, an intended eastward movement may instead result in a northward or southward move, each with probability 0.1. Collisions with obstacles or boundaries cause the agent to remain in its current cell.

*Partial observable multi-agent system.* The underlying environment state  $s_e$  is partially observed by a multi-agent system, which is a directed sensor network. The sensor can change its limited field of view by choosing different angles of viewing. The sensing ranges of these sensors are illustrated in the top two panels of Figure 2. Each sensor has four possible sensing actions: NW, NE, SW, and SE. A sensor may detect the robot if the robot lies within the chosen sensing direction. If the robot is detected, the sensor will receive the exact robot location. We consider a weak sensor scenario. The detection process is stochastic: if the robot is within range, it is detected with probability 0.5; otherwise, with probability 0.5, the sensor receives a null observation ( $n$ ). When the robot is outside the sensing range, the sensor also receives a null observation with probability 1. Note that this multi-agent system, including the uncontrollable environment agent, can be modeled as a Dec-POMDP satisfying Assumption 2.

*Inference objective.* We denote  $Z$  be the estimate of the unknown robot’s type. This corresponds to the case with inferring the environment secret (section 3.3)  $Z = f(S_{e,0:T})$  is defined such that

$f(S_{e,0:T}) = \mathbf{b}$  if  $S_{e,0} = (x, y, \mathbf{b})$  for any initial state  $(x, y)$ . The objective is twofold: 1. select 5 locations for allocating sensors, and 2. compute an optimal sensing policy for each selected sensor such that the collective observations of the sensors for inferring the type of the robot.

*Observation-based policy space.* For each selected sensor, we employ a long short-term memory (LSTM) neural network-based policy that directly processes observation sequences to action decisions. We implement Algorithm 1 in which the individual sensor’s optimal sensing strategy is computed using the policy gradient method in [22]. Figure 3a illustrates the convergence trend in a setting where the dynamics are deterministic and the sensor range is small <sup>2</sup>.

Because maximizing mutual information is equivalent to minimizing conditional entropy, which is often more intuitive to interpret, we present the experimental results in terms of conditional entropy. Initially, when no sensors are allocated, the entropy  $H(Z)$  equals 1, corresponding to the uncertainty in the initial type distribution. After the first iteration, sensor 4 is selected, and its optimal perception policy reduces the conditional entropy (given its observation) to 0.78. In the second iteration, sensor 2 is selected, and the joint observations from sensors 4 and 2 further reduce the conditional entropy to 0.6. We observe that the policy parameter for each selected sensor converges quickly, within the first 100 iterations. This explains the drops shown in Fig. 3a.

This iterative process continues until all five sensors are allocated. Upon convergence, the conditional entropy  $H(Z|Y; \theta^*)$ , where  $\theta^*$  denotes the parameters of the computed optimal decentralized policies, is approximately 0.367. This result indicates that the collective observations provide substantial information about the ground robot’s type on average. The selected sensors are 2, 3, 4, 6, 8—positions that are close to the goals. Notably, sensors located near the robot’s initial position were not selected. We also construct confusion matrices to visualize the classification performance (Figure 3d). These results show that our method enables the sensors to accurately

<sup>2</sup>We sample  $M = 100$  trajectories and set the horizon  $T = 15$  for each iteration. The fixed learning rate of the policy gradient algorithm is set to be 0.001. The hidden dimensions of all layers are set to be 64. We run  $N = 1000$  iterations for each sensor on the 12th Gen Intel(R) Core(TM) i7-12700; the average time consumed for one iteration is 1.3 seconds. The code is available at <https://github.com/AronYoung414/multi-agent-active-perception-grid-world>.

**Table 1: Baseline comparison between IMAS<sup>2</sup> and IPG.**

	Resulting Entropy	Inference accuracy	Time per iteration
IMAS <sup>2</sup>	0.493	86.0%	1.58 s
Fixed Selector (IPG)	0.525	75.5%	7.62 s
Random Selector (IPG)	0.558	70.7%	7.63 s
Visibility-Based Selector (IPG)	0.502	84.1%	7.63 s

infer the robot’s type in both nominal and adversarial settings. The inference accuracy is high as 88%.

We further evaluate our algorithm under different sensing configurations and environment dynamics. Specifically, we consider two setups: selecting three sensors from five candidate locations (sensor 0, 1, 2, 3, 4) and selecting four sensors from seven candidate locations (sensor 0, 1, 2, 3, 4, 5, 6). Both deterministic and stochastic motion models are tested. In addition, we vary the sensing range—small versus large—as illustrated in the bottom two panels of Figure 2.

Figures 3b and 3c summarize the resulting conditional entropies under these configurations. For the small sensing range (Figure 3b), the entropy decreases from about 0.53 to 0.32 in the deterministic setting and from 0.70 to 0.48 in the stochastic setting as the number of sensors increases from 3 to 5. This indicates that adding sensors consistently improves the ability to infer the robot’s type or trajectory. However, the entropy values remain relatively high because of the limited coverage and observation overlap.

When the sensor range is enlarged (Figure 3c), the performance improves significantly. The entropy drops from roughly 0.32 to 0.09 in the deterministic case and from 0.43 to 0.24 in the stochastic case as the number of sensors increases. These results highlight two important trends: 1. deterministic environments yield lower residual uncertainty than stochastic ones, since robot behavior is more predictable; and 2. wider sensing coverage substantially enhances information gain, even with the same number of deployed sensors. Overall, the quantitative results demonstrate that the proposed IMAS<sup>2</sup> algorithm effectively balances agent selection and policy optimization to reduce uncertainty in cooperative perception tasks.

**Baseline Comparison.** Existing approaches cannot directly address the joint problem of sensing-agent selection and decentralized active perception. The sensing approaches in [1, 21] employ Bayesian inference mechanisms similar to those used in IMAS<sup>2</sup>. However, these works focus on *passive* observation models and do not design or optimize *active sensing policies* that control how information is gathered. And the key challenge is that our objective is *mutual information*, rather than a cumulative reward/value function. Consequently, standard value-function-based MARL algorithms (e.g., MADDPG [14], MAPPO [26]) are not directly applicable, and classical Dec-POMDP solvers require an explicit reward structure or a belief-based value function. We therefore compare against a variant of the Independent Policy Gradient (IPG) method [4], which is a decentralized gradient-based method that optimizes each agent’s policy directly via policy gradients without relying on a centralized value function. Since IPG does not include a mechanism for selecting an optimal sensor subset, we provide it with a fixed chosen group of sensors (1, 3, 5, 6, 7), randomly chosen group (8, 5, 4, 6, 1)

of sensors, and visibility-based selector (The set (0, 2, 4, 5, 8) covers the most area) from the nine available candidates. Both methods are evaluated under the stochastic environment with a large sensor range.

As shown in Table 1, the proposed IMAS<sup>2</sup> algorithm achieves a lower conditional entropy (0.493) compared to IPG (0.525, 0.558, 0.502), indicating improved estimation accuracy. Actually, IMAS<sup>2</sup> gives a higher inference accuracy (86.0%) in the test environment compared to IPG (75.5%, 70.7%, 84.1%). Moreover, IMAS<sup>2</sup> converges substantially faster, requiring only 1.5 s per iteration—approximately 5.06 times faster than the IPG baseline (7.6 s). These results demonstrate that IMAS<sup>2</sup> effectively balances computational efficiency and information gain in decentralized active perception tasks.

## 6 CONCLUSION

This paper presented a unified framework for joint agent selection and decentralized policy synthesis in cooperative active perception under the Dec-POMDP setting. By formulating the perception objective in terms of mutual information and conditional entropy, we established that, under mild independence assumptions, the resulting objective is monotone and submodular with respect to the subset of selected agents’ policies. Leveraging this property, we developed the IMAS<sup>2</sup> algorithm, which combines submodular optimization with algorithms for active perception planning. Theoretical analysis showed that under a condition for subsequent maximal marginal gains, the proposed algorithm ensures a tight  $(1 - 1/e)$  approximation guarantee despite the infinite continuous policy space. Our experiments in stochastic and deterministic grid-world environments validated the approach, demonstrating that the method effectively solves agent selection and policy optimization to minimize uncertainty in cooperative perception tasks.

Future research could explore the extension to continuous-state and continuous-action Dec-POMDPs and practical applications such as environment monitoring, intrusion detection, or target tracking. Another direction is to investigate scenarios where the perception agents possess imprecise or uncertain knowledge of the model dynamics, requiring robust or adaptive extensions of the framework. Finally, extending the approach to continuous observation spaces—such as camera images or rich sensory data—would further broaden its applicability to real-world multi-robot and autonomous perception systems.

## ACKNOWLEDGMENTS

Research was sponsored by the Army Research Laboratory under Cooperative Agreement Number W911NF-25-2-0045 and by the Army Research Office under Award Number W911NF-22-1-0166. ARO, as the Federal awarding agency, reserves a royalty-free, nonexclusive and irrevocable right to reproduce, publish, or otherwise use this software for Federal purposes, and to authorize others to do so in accordance with 2 CFR 200.315(b). The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the Army Research Laboratory or the U.S. Government.

## REFERENCES

- [1] Stefano V. Albrecht and Peter Stone. 2017. Reasoning about Hypothetical Agent Behaviours and their Parameters. In *Proceedings of the 16th Conference on Autonomous Agents and MultiAgent Systems (AAMAS '17)*. International Foundation for Autonomous Agents and Multiagent Systems, 547–555.
- [2] Graeme Best, Oliver M Cliff, Timothy Patten, Ramgopal R Mettu, and Robert Fitch. 2019. Dec-MCTS: Decentralized planning for multi-robot active perception. *The International Journal of Robotics Research* 38, 2-3 (2019), 316–337. <https://doi.org/10.1177/0278364918755924> arXiv:<https://doi.org/10.1177/0278364918755924>
- [3] Micah Corah and Nathan Michael. 2021. Scalable Distributed Planning for Multi-Robot, Multi-Target Tracking. *CoRR* abs/2107.08550 (2021). arXiv:2107.08550 <https://arxiv.org/abs/2107.08550>
- [4] Constantinos Daskalakis, Dylan J. Foster, and Noah Golowich. 2020. Independent policy gradient methods for competitive reinforcement learning. In *Proceedings of the 34th International Conference on Neural Information Processing Systems (Vancouver, BC, Canada) (NIPS '20)*. Curran Associates Inc., Red Hook, NY, USA, Article 464, 14 pages.
- [5] Satoru Fujishige. 2005. *Submodular functions and optimization*. Vol. 58. Elsevier.
- [6] Tian He, P. Vicaire, Ting Yan, Liqian Luo, Lin Gu, Gang Zhou, R. Stoleru, Qing Cao, J.A. Stankovic, and T. Abdelzaher. 2006. Achieving Real-Time Target Tracking Using Wireless Sensor Networks. In *12th IEEE Real-Time and Embedded Technology and Applications Symposium (RTAS'06)*. 37–48. <https://doi.org/10.1109/RTAS.2006.9>
- [7] George Kantor, Sanjiv Singh, Ronald Peterson, Daniela Rus, Aavek Das, Vijay Kumar, Guilherme Pereira, and John Spletzer. 2006. *Distributed Search and Rescue with Robot and Sensor Teams*. Springer Berlin Heidelberg, Berlin, Heidelberg, 529–538. [https://doi.org/10.1007/10991459\\_51](https://doi.org/10.1007/10991459_51)
- [8] MHR. Khouzani and Pasquale Malacaria. 2017. Leakage-Minimal Design: Universality, Limitations, and Applications. In *2017 IEEE 30th Computer Security Foundations Symposium (CSF)*. 305–317. <https://doi.org/10.1109/CSF.2017.40>
- [9] Andreas Krause and Carlos Guestrin. 2007. Near-optimal observation selection using submodular functions. In *Proceedings of the 22nd National Conference on Artificial Intelligence - Volume 2 (Vancouver, British Columbia, Canada) (AAAI'07)*. AAAI Press, 1650–1654.
- [10] Chris Kreucher, Keith Kastella, and Alfred O. Hero. 2005. Sensor management using an active sensing approach. *Signal Process.* 85, 3 (March 2005), 607–624. <https://doi.org/10.1016/j.sigpro.2004.11.004>
- [11] Rajiv Ranjan Kumar, Pradeep Varakantham, and Akshat Kumar. 2017. Decentralized planning in stochastic environments with submodular rewards. In *Proceedings of the Thirty-First AAAI Conference on Artificial Intelligence (San Francisco, California, USA) (AAAI'17)*. AAAI Press, 3021–3028.
- [12] Mikko Lauri and Frans A. Oliehoek. 2020. Multi-agent active perception with prediction rewards. In *Proceedings of the 34th International Conference on Neural Information Processing Systems (Vancouver, BC, Canada) (NIPS '20)*. Curran Associates Inc., Red Hook, NY, USA, Article 1145, 11 pages.
- [13] Jing Li, Jing Xu, Fangwei Zhong, Xiangyu Kong, Yu Qiao, and Yizhou Wang. 2020. Assisted Multi-Camera Collaboration for Active Object Tracking. *CoRR* abs/2001.05161 (2020). arXiv:2001.05161 <https://arxiv.org/abs/2001.05161>
- [14] Ryan Lowe, Yi Wu, Aviv Tamar, Jean Harb, Pieter Abbeel, and Igor Mordatch. 2017. Multi-Agent Actor-Critic for Mixed Cooperative-Competitive Environments. *arXiv preprint arXiv:1706.02275* (2017). arXiv:1706.02275 [cs.LG]
- [15] Timothy L. Molloy and Girish N. Nair. 2023. Smoother Entropy for Active State Trajectory Estimation and Obfuscation in POMDPs. *IEEE Trans. Automat. Control* 68, 6 (2023), 3557–3572. <https://doi.org/10.1109/TAC.2023.3250159>
- [16] George L. Nemhauser, Laurence A. Wolsey, and Marshall L. Fisher. 1978. An analysis of approximations for maximizing submodular set functions—I. *Mathematical Programming* 14, 1 (1978), 265–294.
- [17] Ariel Rosenfeld, Noa Agmon, Oleg Maksimov, Amos Azaria, and Sarit Kraus. 2015. Intelligent agent supporting human-multi-robot team collaboration. In *Proceedings of the 24th International Conference on Artificial Intelligence (Buenos Aires, Argentina) (IJCAI'15)*. AAAI Press, 1902–1908.
- [18] Paul E. Rybski, Sascha A. Stoeter, Michael D. Erickson, Maria Gini, Dean F. Hougen, and Nikolaos Papanikolopoulos. 2000. A team of robotic agents for surveillance. In *Proceedings of the Fourth International Conference on Autonomous Agents (Barcelona, Spain) (AGENTS '00)*. Association for Computing Machinery, New York, NY, USA, 9–16. <https://doi.org/10.1145/336595.336607>
- [19] Yash Satsangi, Shimon Whiteson, and Frans Oliehoek. 2015. Exploiting Submodular Value Functions for Faster Dynamic Sensor Selection. *Proceedings of the AAAI Conference on Artificial Intelligence* 29, 1 (Mar. 2015). <https://doi.org/10.1609/aaai.v29i1.9666>
- [20] Yash Satsangi, Shimon Whiteson, Frans A. Oliehoek, and Matthijs T. J. Spaan. 2020. Exploiting Submodular Value Functions For Scaling Up Active Perception. *CoRR* abs/2009.09696 (2020). arXiv:2009.09696 <https://arxiv.org/abs/2009.09696>
- [21] Elnaz Shafipour Yourdshahi, Marcos Antonio do Carmo Alves, Arpit Varma, et al. 2022. On-line estimators for ad-hoc task execution: learning types and parameters of teammates for effective teamwork. *Autonomous Agents and Multi-Agent Systems* 36, 2 (2022), 45.
- [22] Chongyang Shi, Shuo Han, Michael Dorothy, and Jie Fu. 2024. Active Perception With Initial-State Uncertainty: A Policy Gradient Method. *IEEE Control Systems Letters* 8 (2024), 3147–3152. <https://doi.org/10.1109/LCSYS.2024.3513896>
- [23] Amarjeet Singh, Andreas Krause, Carlos Guestrin, and William J. Kaiser. 2009. Efficient informative sensing using multiple robots. *J. Artif. Int. Res.* 34, 1 (April 2009), 707–755.
- [24] Binglu Wang, Lei Zhang, Zhaozhong Wang, Yongqiang Zhao, and Tianfei Zhou. 2023. CORE: Cooperative Reconstruction for Multi-Agent Perception. arXiv:2307.11514 [cs.CV] <https://arxiv.org/abs/2307.11514>
- [25] Jing Xu, Fangwei Zhong, and Yizhou Wang. 2020. Learning multi-agent coordination for enhancing target coverage in directional sensor networks. In *Proceedings of the 34th International Conference on Neural Information Processing Systems (Vancouver, BC, Canada) (NIPS '20)*. Curran Associates Inc., Red Hook, NY, USA, Article 843, 12 pages.
- [26] Chao Yu, Akash Velu, Eugene Vinitzky, Jiaxuan Gao, Yu Wang, Alexandre Bayen, and Yi Wu. 2021. The Surprising Effectiveness of PPO in Cooperative Multi-Agent Games. In *Advances in Neural Information Processing Systems*, Vol. 34. 24611–24624.