

Multi-agent Task-Driven Exploration via Intelligent Map Compression and Sharing

Evangelos Psomiadis¹, Dipankar Maity², and Panagiotis Tsiotras¹

¹ D. Guggenheim School of Aerospace Engineering, Georgia Institute of Technology,
Atlanta, GA, 30332-0150, USA,
{`epsomiadis3`, `tsiotras`}@gatech.edu

² Department of Electrical and Computer Engineering, University of North Carolina
at Charlotte, NC, 28223-0001, USA,
`dmaity@charlotte.edu` *

Abstract. This paper investigates the task-driven exploration of unknown environments with mobile sensors communicating compressed measurements. The sensors explore the area and transmit their compressed data to another robot, assisting it to reach its goal location. We propose a novel communication framework and a tractable multi-agent exploration algorithm to select the sensors' actions. The algorithm uses a task-driven measure of uncertainty, resulting from map compression, as a reward function. We validate the efficacy of our algorithm through numerical simulations conducted on a realistic map and compare it with alternative approaches. The results indicate that the proposed algorithm effectively decreases the time required for the robot to reach its target without causing excessive load on the communication network.

1 Introduction

In recent years, there has been a notable rise in the use of autonomous robot teams undertaking intricate tasks [1–3]. During these operations, effective communication between the robots becomes crucial in order to optimize the overall performance of the team. To avoid unnecessary delays, it is essential for the robots to be cognizant of the network's bandwidth limitations and incorporate them in their control loop [4], by compressing their measurements accordingly.

Resource-aware decision-making is indispensable for achieving assured autonomy in challenging conditions involving multi-agent interactions under severe communication and computational constraints [5–10]. In particular, the coupled compression and control problem has gathered significant attention within both the classical control [11–14] and robotics [15–18] communities. Although the problem of designing the optimal quantizer, even for a single agent with a single task, has been proven to be intractable [19, 20], alternative approaches have been proposed that tackle the problem by selecting the optimal quantizer from a given set of quantizers [21, 22]. In our previous work [15], we adopted a similar approach for mobile robot path-planning using compressed measurements of

* The work was supported by the ARL grant DCIST CRA W911NF-17-2-0181.




Fig. 1: (a) Finest resolution grid; (b) Compressed grid; (c) Variance of the compressed grid, with darker cells indicating higher variance.

mobile sensors having predefined paths. In this work, we develop a compression framework similar to [15] for a set of mobile sensors to intelligently compress their data in a task-driven manner. We also consider an exploration strategy for these mobile sensors to uncover unknown regions in a task-driven manner.

Distributed sensing or sensor placement has been widely studied in the estimation and controls communities [23]. Different performance metrics have been used to drive optimal sensor placement such as entropy, mutual information, Kullback–Leibler divergence [24], energy-constraints [25], and value-of-information [26]. Although most of the literature includes stationary sensor placement, there have been notable extensions to mobile sensors as well [27, 28].

Inspired by distributed sensing, many researchers have studied the problem of robot exploration and motion-planning to minimize the uncertainty of mapping and perception. Most information-theoretic exploration approaches utilize Shannon’s entropy as the information metric [29, 30]. In [31] the authors proposed the Cauchy-Schwarz Quadratic Mutual Information as an alternative to Shannon’s entropy for faster computation. Unlike the existing literature, our focus is on minimizing the uncertainty of compression in the environment.

Contributions: Our contribution consists of three main elements. Firstly, we expand the conventional grid-world compression scheme by assigning not only a single value per compressed cell but a pair of values, representing both the mean and variance of the cell. Secondly, this extension enables us to approximate the compression uncertainty and improve the estimation process, in comparison to previous frameworks relying solely on the mean. Finally, we use the previous results to propose a tractable multi-agent task-driven exploration algorithm to assist another robot in reaching a target by reducing the compression uncertainty of the environment during the exploration phase.

2 Preliminaries: Grid World and Abstractions

We adopt the premise that the environment is represented by a 2D static grid denoted by $\mathcal{M} \subset \mathbb{R}^2$ (e.g., occupancy grid as proposed in [32, 33], elevation map, or a more general discretized cost map), with N being the total number of cells in




Fig. 2: Example of a set of abstractions. The grey area denotes no information.

$\mathcal{M}$ and $\mathbf{p}_j$ being the position (e.g., the center of the cell) of the j -th cell in $\mathcal{M}$. We define the index set of the finest resolution cells within $\mathcal{M}$ as $M = \{j : \mathbf{p}_j \in \mathcal{M}\}$. A robot is equipped with onboard sensors capable of observing a 2D grid of size $w \times h \leq N$ as shown in Fig. 1(a). The vector $x \in [0, 1]^N$ contains the values of the grid cells, where the value of the j -th component of x , denoted by x_j , lies in the range $[0, 1]$ for all $j = 1, \dots, N$. Here x_j represents the traversability of the cell, where $x_j = 1$ indicates a non-traversable cell and $x_j = 0$ denotes an obstacle-free cell. All robots have the same understanding of cell traversability.

We define as an *abstraction* a multi-resolution compressed representation of the grid $\mathcal{M}$, as in Fig. 1(b). Each abstraction is characterized by its *compression template*, indicating how the grid cells are compressed. The structure of an abstraction is based on the robot's sensed map and is manually crafted beforehand. Fig. 2 shows examples of a few possible abstractions for a 7×7 grid. Various techniques have been used to determine the value of a compressed cell [6, 34, 35]. In our previous work [15], we considered the value of the compressed cell to be the average of the values of the underlying cells. This way, each abstraction can be thought of as a linear mapping $\mathcal{A}^\theta : [0, 1]^N \rightarrow [0, 1]^{k^\theta}$, where $k^\theta \leq wh \leq N$ is the number of cells in the compressed representation utilizing abstraction θ for a given robot's local sensed map. Specifically, for a high resolution map $x \in [0, 1]^N$, the values of the cells at abstraction θ are given by $o = \mathcal{A}^\theta x$, where $\mathcal{A}^\theta \in \mathbb{R}^{k^\theta \times N}$. For example, the dimensions in Fig. 2 are $w = 7$, $h = 7$, and the number of compressed cells for abstractions $\theta = 8$ and 9 is $k^\theta = 7$. We will denote the j -th finest resolution cell as $j \in M$ and the i -th compressed cell as $i \in I^\theta$, where I^θ is the index set of compressed cells obtained from $\mathcal{M}$ through abstraction θ having cardinality equal to k^θ . With a slight abuse of notation, we will use the notation $j \in i$ to denote that the finest resolution cell in $\mathcal{M}$ with index j belongs to the compressed cell in I^θ having index i .

In this study, each i -th cell in I^θ is assigned a pair of values, the mean and the variance. The communication of the variance has a two-fold impact. Firstly, it serves as an uncertainty indicator due to the compression. Secondly, it provides better estimates along with the mean. We define the variance V_i of the i -th cell of abstraction θ as the normalized squared deviation from the mean, that is,

$$V_i = \sum_{j \in M} \mathcal{A}_{ij}^\theta (x_j - o_i)^2, \quad i \in I^\theta. \quad (1)$$

Note that $\mathcal{A}_{ij}^\theta = 1/N_i^\theta$ if $j \in i$, and $\mathcal{A}_{ij}^\theta = 0$ otherwise, where N_i^θ is the number of non-zero elements in the i -th row of $\mathcal{A}^\theta$. Fig. 1(c) shows the variance of the compressed grid presented in Fig. 1(b). Darker cells indicate higher variance.

2.1 Communication of Abstracted Environments

Consider a pair of mobile robots, named as Actor and Sensor, that communicate with each other. At each timestep t , the Sensor chooses an abstraction $\theta_t \in \Theta$ to compress its local map and transmits it to the Actor, where $\Theta = \{1, \dots, K\}$ is the set of available abstractions (i.e., codebook) the two robots have agreed upon (e.g., see Fig. 2 for a collection of templates). Specifically, the Sensor transmits the abstraction triplet (o_t, V_t, θ_t) , along with its current position $\mathbf{p}_t^s$, where V_t is the compression variance given in (1), and $o_t = \mathcal{A}^{\theta_t} x$ with x containing the measurements of the Sensor. Having received the transmitted θ_t and position $\mathbf{p}_t^s$, the Actor computes $\mathcal{A}^{\theta_t}$ and attempts to reconstruct x from o_t and V_t .

We denote by n_m the number of bits required to transmit the compressed value o_t and variance V_t , and n_a the number of bits required to transmit an abstraction index θ . Hence, the total number of bits n^θ required to transmit the abstraction triplet employing abstraction θ is:

$$n^\theta = k^\theta n_m + n_a. \quad (2)$$

3 Problem Formulation

Consider a team of mobile robots with different objectives navigating through an unfamiliar environment $\mathcal{M} \subset \mathbb{R}^2$ cluttered with static obstacles. Each robot has its own sensing equipment capable of observing a portion of the environment (local map) as it moves. The team consists of the Actor robot and multiple mobile Sensors. The Actor's goal is to reach a specified target location in minimum time. Meanwhile, the Sensors aim to support the Actor in achieving its objective by communicating informative abstractions of their local maps to the Actor at each timestep. The Sensors can be regarded, for example, as surveillance drones capable of moving over ground obstacles, while the Actor is assumed to be a ground robot, that has to circumnavigate these obstacles.

Let $\mathbf{p}_t^A, \mathbf{p}_t^s \in \mathcal{M}$ be the Actor and each Sensor's position, respectively, at timestep t , where s denotes the Sensor index, and let $u_t^A \in U^A, u_t^s \in U^S$ denote the control actions of the Actor and each Sensor at time t selected from a finite set of control actions U^A and U^S , respectively.

3.1 Problem Statement

Given the uncertainty of cell values caused by compression and the initial positions of the Actor $\mathbf{p}_0^A$ and the Sensors $\mathbf{p}_0^s$, as well as the Actor's destination $\mathbf{p}_G^A$ and the set of abstractions Θ , we wish to design an exploration strategy for the Sensors to minimize compression uncertainty near the Actor's path. Each Sensor communicates to the Actor the triplet $(o_t^s, V_t^s, \theta_t^s)$, as detailed in Section 2.1, along with its current position $\mathbf{p}_t^s$. The Actor utilizes the variance of each Sensor V_t^s to estimate the map compression uncertainty and selects each Sensor's action u_t^s to reduce uncertainty near its path. Our primary goal is for the Actor to reach its destination in minimum time (shortest path) by using the accumulated measurements to compute its control u_t^A .

4 Compression Uncertainty

In this section, we quantify the map uncertainty resulting from the unknown environment, and the history of transmitted abstractions $\theta_{0:t}$ until time t . The quantified result is represented by a vector $h_t \in \mathbb{R}^N$, which approximates the uncertainty of each cell in $\mathcal{M}$. The Actor employs h_t to compute the trajectories/actions of the Sensors, as elaborated later in Section 5.3.

First, we rewrite (1) as follows

$$x^\top [\text{diag } \mathcal{A}_i^\theta] x = o_i^2 + V_i, \quad (3)$$

where $[\text{diag } \mathcal{A}_i^\theta] \in [0, 1]^{N \times N}$ is a diagonal matrix with the i -th row of $\mathcal{A}^\theta$ as its diagonal.

Equation (3) along with $o = \mathcal{A}^\theta x$ describes a set in $\mathbb{R}^N$ that serves as the estimation domain imposed by abstraction θ . To make the problem computationally tractable, we relax (3) by replacing the equality with an inequality. Hence, a valid estimate $\hat{x}$ of $x \in [0, 1]^N$ belongs in the set

$$C^\theta = \left\{ x \in \mathbb{R}^N : \begin{array}{ll} x^\top [\text{diag } \mathcal{A}_i^\theta] x \leq o_i^2 + V_i, & i = 1, 2, \dots, k^\theta \\ \mathcal{A}^\theta x = o & \\ 0 \leq x_j \leq 1, & j = 1, \dots, N \end{array} \right\}. \quad (4)$$

Note that relaxing (3) makes C^θ convex, which significantly improves the computational efficiency in the subsequent analysis (e.g., see Section 5.1).

Each robot engages in the estimation process individually. Each Sensor uses only its own abstractions to estimate x , while the Actor utilizes its own measurements and the communicated abstractions from all the Sensors. Considering the history of abstractions $\theta_{0:t}^s$ of Sensor s up to time t , the Sensor's measurements are described by the set $C_t^s = \cap_{\tau=0}^t C_{\tau}^{\theta_\tau^s}$. The Actor's measurement of the j -th cell in $\mathcal{M}$ can be described by $\alpha_j x = o_j$, where $\alpha_j \in \mathbb{R}^{1 \times N}$ has zeros everywhere except at the j -th element. Let $C_t^{A'}$ be the set containing all of the Actor's measurements up to time t . The Actor can then use these accumulated measurements to form a smaller set, defined as $C_t^A = (\cap_{s \in S} C_t^s) \cap C_t^{A'}$, which it utilizes for its estimation process.

We define the compression uncertainty of the j -th cell in $\mathcal{M}$ as

$$\mathcal{H}_j = \max_{x, \hat{x} \in C_t^A} |x_j - \hat{x}_j|. \quad (5)$$

However, as the j -th cell may be included in multiple compressed cells, computing $\mathcal{H}_j$ requires considering all the compressed cells within which it is contained. To make the computation of $\mathcal{H}_j$ tractable, we derive an upper bound.

Proposition 1. *Let x_j be the true value of the j -th cell and $\hat{x}_j$ be a valid estimate of it. Let $I_j = \{i : i \in I^{\theta_{0:t}} \text{ and } j \in i\}$. Then,*

$$\mathcal{H}_j \leq h_j, \quad (6a)$$

$$h_j = \min_{i \in I_j} \{\nu_i + o_i\} + \min_{i \in I_j} \{\nu_i - o_i\}, \quad (6b)$$

where $\nu_i = \sqrt{N_i^\theta V_i}$.




Fig. 3: Flowchart of the proposed framework’s architecture at timestep t . The Actor shares its current optimal path π_t^* to the Sensors along with their actions $u_{s,t}$. The Sensors utilize π_t^* to select the optimal abstraction of their respective local maps $\theta_{s,t}^* \in \Theta$ and send it to the Actor to aid in reaching its target.

Proof. From Equation (1)

$$|x_j - o_i| \leq \nu_i, \quad \forall i \in I_j. \quad (7)$$

It follows that

$$\max_{i \in I_j} \{o_i - \nu_i\} \leq x_j \leq \min_{i \in I_j} \{o_i + \nu_i\}. \quad (8)$$

A valid estimate $\hat{x}_j$ of x_j must also satisfy

$$\max_{i \in I_j} \{o_i - \nu_i\} \leq \hat{x}_j \leq \min_{i \in I_j} \{o_i + \nu_i\}, \quad (9)$$

since the true cell value lies in this range. Combining (8) and (9), we obtain the desired result.

5 Framework Architecture

In this section, we introduce the proposed framework, depicted in Fig. 3. The framework leverages the same Decoder and Path Planner as in [15].

5.1 Decoder

The decoder (see Fig. 3) generates estimates for $x \in [0, 1]^N$, representing the cell values, by merging the accumulated measurements. Each robot utilizes its own constraint set C_t to compute its estimates (C_t^A for Actor and C_t^s for Sensor s). Hence, the estimation vector is given by solving the convex optimization problem

$$\hat{x}_t^* = \arg \min_{\hat{x}_t \in C_t} \|\hat{x}_t - \frac{1}{2}\mathbf{1}\|^2, \quad (10)$$

where we assume that x follows an unknown distribution such that, for all $j \in M$, $\mathbb{E}[x_j] = 0.5$. The intuition for this choice and the proof can be found in [15].




Fig. 4: (a) Example of a discretized environment. **S** denotes the starting cell while **G** denotes the target; (b) Associated graph of the environment employing the set of actions $U = \{\text{UP, DOWN, LEFT, RIGHT}\}$, along with the optimal path (blue nodes) with cost function (11); (c) Path weights computed using (13).

5.2 Path Planner

Consider the graph $\mathcal{G} = (\mathcal{N}, \mathcal{E})$ associated with the discretized map $\mathcal{M}$, where $\mathcal{N}$ represents the set of vertices and $\mathcal{E}$ the set of edges. With a slight abuse of notation, we denote both the cell positions in $\mathcal{M}$ and the graph vertices $\mathcal{N}$ as $\mathbf{p}$, since they coincide. An edge connects two vertices if the Actor can move between them using a control action $u_A \in U^A$. We assume that the traversal time of a cell is proportional to its value (difficulty of traversal) plus a constant (penalty for movement). The cost of traversing a vertex is given by [15, 36]

$$c_\epsilon(\mathbf{p}) = \begin{cases} x_{\mathbf{p}} + a, & \text{if } \mathbf{p} \in P_\epsilon, \\ N(\epsilon + a), & \text{otherwise,} \end{cases} \quad (11)$$

where $x_{\mathbf{p}} \in [0, 1]$ is the value of the cell at position $\mathbf{p}$, a is a constant penalty for movement, and $P_\epsilon = \{\mathbf{p} \in \mathcal{N} : x_{\mathbf{p}} \leq \epsilon\}$ is the set of cells meeting a feasibility condition, with $\epsilon \in [0, 1]$ being the threshold for considering a cell to be an obstacle.

Utilizing (11), the Actor's optimal path is given by

$$\pi^* = \arg \min_{\pi \in \Pi} \sum_{\mathbf{p} \in \pi} c_\epsilon(\mathbf{p}), \quad (12)$$

where Π is the set of paths with the first element being the Actor's position at timestep t , denoted by $\mathbf{p}_t^A$, and the last element being its goal location $\mathbf{p}_G^A$.

Fig. 4(a) presents an occupancy grid with obstacles, while Fig. 4(b) illustrates the associated graph with the Actor's optimal path (computed, for instance, using Dijkstra's algorithm [37]).

5.3 Sensor Action Selector

The actions of the Sensors are selected to minimize the Actor's compression uncertainty. To accomplish this, the Sensors' actions are selected based on a

task-driven metric for uncertainty, denoted by $R_t \in \mathbb{R}^N$. This metric relies on the Actor's current path π_t^* and the uncertainty of the whole map. In this study, the Sensors' actions are selected by the Actor in a centralized manner (see Fig. 3). However, in the case of a single Sensor, the Sensor can select its own actions, as there are no abstractions from other Sensors that cause compression uncertainty.

To compute R_t , first we construct an area of interest around π_t^* by assigning weights to each cell within M . Given that the Actor will likely explore the vicinity of its current position in subsequent timesteps, and that its path may vary in the future, it is more advantageous for the Sensors to prioritize exploration in the vicinity of the Actor's target. Hence, the weight of each cell is given by

$$w_t(\mathbf{p}) = \frac{|\pi_d|}{|\pi_t^*|} e^{-\frac{d^2}{2v}}, \quad (13a)$$

$$d = \min_{\mathbf{p}_\pi \in \pi_t^*} \|\mathbf{p} - \mathbf{p}_\pi\|, \quad (13b)$$

where $\mathbf{p} \in \mathcal{M}$, v is a parameter characterizing the width of the region of interest around the path, $|\pi_t^*|$ is the total number of nodes (length) in π_t^* , and $|\pi_d|$ is the length of the portion of π_t^* , starting from the current location $\mathbf{p}_{A,t}$ up to the argument of minimization of (13b). Fig. 4(c) depicts these path weights.

Hence, we define the task-driven metric of uncertainty R_t as follows

$$R_t = W_t \circ h_t, \quad (14)$$

where $W_t \in \mathbb{R}^N$ contains the weights $w_t(\mathbf{p})$ of every cell computed by (13), $\circ$ is the Hadamard product, and $h_t \in \mathbb{R}^N$ contains the uncertainty given by (6).

In this study, the action of each Sensor is computed by solving an infinite-horizon discounted Markov Decision Process (MDP) $\langle \mathcal{L}^s, U^s, \tilde{R}_t^s, \gamma, p \rangle$ at each timestep t . Here, $\mathcal{L}^s$ represents the set of states, where each state corresponds to the local map of Sensor s , denoted by L^s , and U^s is the set of actions. To provide a tractable and computationally efficient solution, we restrict the states to a neighborhood around the Sensor's current local map. Let ℓ be the radius of that neighborhood and let $N' \leq N$ be the number of states in $\mathcal{L}^s$. The parameter $\gamma \in (0, 1)$ is the discount factor, and $p : \mathcal{L}^s \times \mathcal{L}^s \times U^s \rightarrow \{0, 1\}$ is the transition function with transition matrix $P^s \in \mathbb{R}^{N' \times N'}$.

Each element of the reward vector $\tilde{R}_t^s \in \mathbb{R}^{N'}$ is defined as the sum of rewards obtained at each position $\mathbf{p}$ within the Sensor state L^s

$$\tilde{R}_t^s(L^s) = \sum_{\mathbf{p} \in L^s} R_t(\mathbf{p}). \quad (15)$$

The value function is defined as the vector

$$\mathcal{V}_t^s(\mu^s) = \sum_{\tau=0}^{\infty} \gamma^\tau (P^s(\mu^s))^\tau \tilde{R}_t^s, \quad (16)$$

where μ^s denotes the policy that selects the action at each state L^s . We obtain the optimal policy at each state L^s by $\mu^{s*}(L^s) = \arg \max_{\mu^s} \mathcal{V}_t^s(\mu^s)|_{L^s}$. The Sensor's action at each timestep t is therefore given by

$$u_t^s = \mu^{s*}(L_t^s), \quad (17)$$

Algorithm 1 The Sensor Action Selector's Algorithm

Input: $\pi_t^*, \mathcal{A}_{0:t}, o_{0:t}, V_{0:t}, \mathbf{p}^S, \ell, U^S, \gamma, P^s$
Output: u_t^S

- 1: $h_t \leftarrow \text{UNCERTAINTY}(\mathcal{A}_{0:t}, o_{0:t}, V_{0:t})$ using (6)
- 2: $W_t \leftarrow \text{WEIGHTS}(\pi_t^*)$ using (13)
- 3: $R_t \leftarrow \text{TASKDRIVENUNCERTAINTY}(h_t, W_t)$ using (14)
- 4: $S_{prev} \leftarrow \emptyset$ ▷ Initialize Sensor counting
- 5: **for all** $s \in S$ **do**
- 6: $\mathcal{L}^s \leftarrow \text{STATES}(\mathbf{p}^s, \ell)$
- 7: $\tilde{R}_t^s(L^s) \leftarrow \text{REWARD}(L^s, R_t) \quad \forall L^s \in \mathcal{L}^s$ using (15)
- 8: $\tilde{R}_t^s \leftarrow \text{OVERLAPPING}(\tilde{R}_t^s, \mathcal{L}^{s \in S_{prev}})$
- 9: $u_t^s \leftarrow \text{VALUEITERATION}(\mathcal{L}^s, U^S, \tilde{R}_t^s, \gamma, P^s)$
- 10: $S_{prev} \leftarrow S_{prev} \cup \{s\}$
- 11: **end for**
- 12: **return** u_t^S

where L_t^s is the local map of Sensor s when at position $\mathbf{p}_t^s$.

The algorithm for selecting the Sensors' actions is outlined in Algorithm 1. Lines 5-11 iteratively compute the action for each Sensor by solving the individual MDPs. Note that Line 8 sets the rewards to zero for the states already considered by the Sensors whose actions were determined in previous iterations, thus preventing overlap. Line 9 performs Value Iteration to solve the MDP, which is the most computationally intensive step in the algorithm.

An alternative approach to obtain the Sensors' actions is to select those with the highest immediate reward, thereby avoiding the need to solve an MDP. We call this method the greedy algorithm and we present its results, along with the proposed method, in Section 6.

5.4 Path Converter

The Actor sends its current optimal path π_t^* obtained from (12) to the Sensors. Then, each Sensor utilizes this information to guide its abstraction selection process separately. This is achieved by constructing W_t using (13).

5.5 Encoder

The encoder of the Sensor selects the optimal abstraction from the given set Θ , to transmit to the Actor. Let $\theta^s \in \Theta$ denote Sensor's s abstraction. For the rest of this section, we will omit the superscript s for simplicity. Given the past selection $\theta_{0:t-1}$, the optimal abstraction at t is derived through the minimization of the following criterion by performing an exhaustive search

$$J_t(\theta_t) = \beta D_t + (1 - \beta) H_t + \lambda(\theta_t), \quad (18a)$$

$$D_t = \|W_t \circ (\tilde{x}_t - \hat{x}_t^*(\theta_{0:t}))\|^2, \quad H_t = \|W_t \circ h_t(\theta_{0:t})\|^2, \quad (18b)$$

$$\theta_t^* = \arg \min_{\theta_t \in \Theta} J_t(\theta_t), \quad (18c)$$

where $\beta \in [0, 1]$ is a trade-off factor, $\tilde{x}_t \in \mathbb{R}^N$ is the sensed cell values by each Sensor until time t , $\hat{x}_t(\theta_{0:t}) \in \mathbb{R}^N$ is the estimation vector computed from (10), $h_t(\theta_{0:t})$ is the uncertainty vector given by (6), and $\lambda(\theta_t)$ is the communication cost (related to n^θ in (2)). It is important for each Sensor to select the optimal abstraction for the Actor's specific decoder (see Fig. 3). Hence, the estimation vector $\hat{x}_t(\theta_{0:t})$ is computed using (10) with the constraint set C_t^s .

6 Experiments

Two different scenarios were studied in a 64×64 environment (see Fig. 5(a)). In the first scenario, we varied the radius of the MDP neighborhood ℓ (see Section 5.3) by conducting simulations with three Sensors. We also present the results of the greedy algorithm. In the second scenario, we compared our proposed algorithm to the algorithm introduced in [15]. This scenario involved a single Sensor. In both scenarios, we performed 100 simulations with varied initial positions of the Sensors. All the robots started moving simultaneously, traversing one cell per timestep. The Sensors have the ability to move over obstacles (e.g., surveillance drones).

The Actor's sensing region is a 5×5 grid around its position, while the Sensors have a field of view of 7×7 . The Actor's initial position is $\mathbf{p}_0^A = (6, 29)$ and its target position is $\mathbf{p}_G^A = (43, 25)$. The Actor's Path Planner parameters in (11) are $a = 0.025$ and $\epsilon = 0.501$, the parameter in (13) is $v = 3.33$, and the discount factor in (16) is $\gamma = 0.9$. The Sensors' encoders utilize a finite set of 10 abstractions, as shown in Fig. 2, where the parameters in (18) are $\beta = 0.9$ and $\lambda(\theta) = 0.05k^\theta$, where the cost of abstraction is proportional to the amount of bits required for transmitting the abstracted map.

6.1 Performance Metrics

We assumed that the time needed to traverse a cell for the Actor is proportional to that cell's value (see Section 5.2). Let π_f include the cells that the Actor traversed to reach its goal, without excluding duplicate cells. Then, the accumulated cost is $\mathcal{C} = \sum_{\mathbf{p} \in \pi_f} c_\epsilon(\mathbf{p})$ and the average cost ratio is

$$r_{\text{cost}} = \frac{1}{n_{\text{sim}}} \sum_{i=0}^{n_{\text{sim}}} \frac{\mathcal{C}(i)}{\mathcal{C}_{\text{max}}}, \quad (19)$$

where $n_{\text{sim}} = 100$ is the total simulation number, $\mathcal{C}(i)$ denotes the accumulated cost for simulation i , and $\mathcal{C}_{\text{max}}$ refers to the accumulated cost in an Uninformed framework, where the Actor reaches its target without any assistance from the Sensors. We also compute the average ratio of the transmitted bits

$$r_{\text{bits}} = \frac{1}{n_{\text{sim}}} \sum_{i=0}^{n_{\text{sim}}} \frac{\mathcal{B}(i)}{\mathcal{B}_{\text{max}}(i)}, \quad (20)$$




Fig. 5: (a) Full Map; (b-c) Simulation snapshots of the proposed exploration framework at timestep $t = 8$ (b) and $t = 61$ (c) with $l = 10$, $\mathbf{p}_{1,0} = (14, 24)$, $\mathbf{p}_{2,0} = (6, 21)$, $\mathbf{p}_{3,0} = (22, 21)$. The Actor's path until t is presented in blue, its predicted path in green, while the Sensors' trajectories in red.

Table 1: Simulation results for Scenario 1

	Greedy	$\ell = 1$	$\ell = 2$	$\ell = 10$
r_{cost}	0.31	0.30	0.30	0.31
r_{bits}	0.98	0.98	1	0.98

where $\mathcal{B} = \sum_{t=0}^T n_t$ and n_t is given in (2) and denotes the number of bits sent by each Sensor at t . If the variance is not transmitted, n_m is halved. We set $n_m = 24$, and $n_a = 4$. The number of the Sensors' steps T is fixed at 60 for all simulations. $\mathcal{B}_{\max}$ refers to the approach that results in the maximum value of the sum $\sum_{i=0}^{n_{\text{sim}}} \mathcal{B}(i)$ for each scenario.

6.2 Multiple Sensors

In the first scenario, we tested our exploration algorithm with a team of three Sensors and different values of the MDP neighborhood parameter ℓ . We also present the results of the greedy algorithm, where the Sensors' actions are selected based on the maximum immediate reward. In each simulation, we varied the initial positions of the Sensors.

The resulting average cost and bit ratios are presented in Table 1. We observe that all the algorithms perform significantly better than the Uninformed framework, reducing the cost accumulated by the Actor in reaching its target by 70%. The algorithm with $\ell = 1$ is the best approach for this scenario, achieving the greatest reduction in accumulated cost with the least amount of communication.

However, it is important to note that setting $\ell = 10$ increases the exploration coverage of the map, since ℓ determines the maximum distance between the Sensors (as indicated by Line 8 in Algorithm 1), which is particularly beneficial for larger maps and numerous Sensors. This can be observed in Figs. 5(b)–(c) where a simulation with $\ell = 10$ is shown. Note that Fig. 5(b) is the timeframe before the Actor corrects its computed path.

Table 2: Simulation results for Scenario 2

	Task-Driven Exploration	Predefined path AS [15]	Predefined path FI [15]
r_{cost}	0.72	0.87	0.67
r_{bits}	0.78	0.16	1

6.3 Single Sensor - Comparison with [15]

We tested our proposed exploration algorithm with a single Actor-Sensor pair and an MDP parameter $\ell = 1$, varying the Sensor’s initial position. We compared our algorithm with the predefined path Abstraction-Selection (AS) framework presented in [15]. We simulated two different predefined paths and took the average of them to compute the cost $\mathcal{C}$ and bits $\mathcal{B}$ of the simulations. The paths followed a square trajectory (as in [15]) with one path traversing the square clockwise and the other counterclockwise. We also present the results of the Fully-Informed (FI) framework from [15], where the Sensor follows a predefined path and transmits all observed cells. In the proposed framework, in order to ensure a fair comparison with the other framework and given that we are employing only one Sensor, we avoided the transmission of variance and relocated the Sensor Action Selector (see Fig. 3) from the Actor to the Sensor.

The resulting average cost ratio and bit ratio are presented in Table 2. The results indicate that the proposed framework performs similarly to the FI framework while significantly reducing communication. Additionally, we find that the predefined AS framework requires the least communication. This is because, in most simulations, the predefined path leads the Sensor to areas that are not of interest, whereas our proposed framework determines the Sensor’s actions by directly considering the area of interest.

7 Conclusions

In this paper, we studied the problem of exploring unknown environments with mobile sensors that communicate compressed measurements. The sensors are deployed to facilitate mobile robot path-planning. We introduce a novel communication framework along with a tractable multi-agent exploration algorithm aimed at optimizing the actions of the sensors. Our algorithm leverages a task-driven measure of compression uncertainty as a reward function. Simulation results demonstrate the effectiveness of our framework in reducing the time required for the robot to reach its target without excessively burdening the communication network. Future extensions will be focused on refining the proposed framework by incorporating a more sophisticated abstraction selection search method to decrease the computational time of the Sensors, as well as exploring extensions to dynamic environments.

References

1. J. Thangavelautham, A. Chandra, E. Jensen, Autonomous robot teams for lunar mining base construction and operation, in *IEEE Aerospace Conference* (2020), pp. 1–16
2. O. Salzman, R. Stern, Research challenges and opportunities in multi-agent path finding and multi-agent pickup and delivery problems, in *19th International Conference on Autonomous Agents and MultiAgent Systems* (Auckland, New Zealand, 2020), pp. 1711–1715
3. J.P. Queralta, J. Taipalmaa, B. Can Pullinen, V.K. Sarker, T. Nguyen Gia, H. Tenhunen, M. Gabbouj, J. Raitoharju, T. Westerlund, Collaborative multi-robot search and rescue: Planning, coordination, perception, and active vision. *IEEE Access* **8**, 191617–191643 (2020)
4. J. Gielis, A. Shankar, A. Prorok, A critical review of communications in multi-robot systems. *Current Robotics Reports* **3**(4), 213–225 (2022)
5. Z. Xu, V. Tzoumas, Resource-aware distributed submodular maximization: A paradigm for multi-robot decision-making, in *IEEE 61st Conference on Decision and Control (CDC)* (Cancun, Mexico, 2022), pp. 5959–5966
6. D.T. Larsson, D. Maity, P. Tsiotras, Q-tree search: An information-theoretic approach toward hierarchical abstractions for agents with computational limitations. *IEEE Transactions on Robotics* **36**(6), 1669–1685 (2020)
7. C. Hireche, C. Dezan, J.P. Diguët, L. Mejias, BFM: A scalable and resource-aware method for adaptive mission planning of UAVs, in *IEEE International Conference on Robotics and Automation (ICRA)* (Brisbane, QLD, Australia, 2018), pp. 6702–6707
8. M.H. Mamduhi, D. Maity, K.H. Johansson, J. Lygeros, Regret-optimal cross-layer co-design in networked control systems—part I: General case. *IEEE Communications Letters* **27**(11), 2874–2878 (2023)
9. D. Maity, M.H. Mamduhi, J. Lygeros, K.H. Johansson, Regret-optimal cross-layer co-design in networked control systems—part II: Gauss-Markov case. *IEEE Communications Letters* **27**(11), 2879–2883 (2023)
10. M.H. Mamduhi, D. Maity, S. Hirche, J.S. Baras, K.H. Johansson, Delay-sensitive joint optimal control and resource management in multiloop networked control systems. *IEEE Transactions on Control of Network Systems* **8**(3), 1093–1106 (2021)
11. D.F. Delchamps, Stabilizing a linear system with quantized state feedback. *IEEE Transactions on Automatic Control* **35**(8), 916–924 (1990)
12. R.W. Brockett, D. Liberzon, Quantized feedback stabilization of linear systems. *IEEE Transactions on Automatic Control* **45**(7), 1279–1289 (2000)
13. G.N. Nair, R.J. Evans, Stabilizability of stochastic linear systems with finite feedback data rates. *SIAM Journal on Control and Optimization* **43**(2), 413–436 (2004)
14. V. Kostina, B. Hassibi, Rate-cost tradeoffs in control. *IEEE Transactions on Automatic Control* **64**(11), 4525–4540 (2019)
15. E. Psomiadis, D. Maity, P. Tsiotras, Communication-aware map compression for online path-planning, in *IEEE International Conference on Robotics and Automation (ICRA)* (Yokohama, Japan, 2024), pp. 12368–12374
16. R. Marcotte, X. Wang, D. Mehta, E. Olson, Optimizing multi-robot communication under bandwidth constraints. *Autonomous Robots* **44**(1), 43–55 (2020)
17. V. Unhelkar, J. Shah, Contact: Deciding to communicate during time-critical collaborative tasks in unknown, deterministic domains, in *AAAI Conference on Artificial Intelligence* (Phoenix, AZ, 2016)

18. G. Damigos, N. Stathouloupoulos, A. Koval, T. Lindgren, G. Nikolakopoulos, Communication-aware control of large data transmissions via centralized cognition and 5G networks for multi-robot map merging. *Journal of Intelligent & Robotic Systems* **110**(22) (2024)
19. M. Fu, Lack of separation principle for quantized linear quadratic gaussian control. *IEEE Transactions on Automatic Control* **57**(9), 2385–2390 (2012)
20. S. Yüksel, A note on the separation of optimal quantization and control policies in networked control. *SIAM Journal on Control and Optimization* **57**(1), 773–782 (2019)
21. D. Maity, P. Tsiotras, Optimal controller synthesis and dynamic quantizer switching for linear-quadratic-Gaussian systems. *IEEE Transactions on Automatic Control* **67**(1), 382–389 (2022)
22. D. Maity, P. Tsiotras, Optimal quantizer scheduling and controller synthesis for partially observable linear systems. *SIAM Journal on Control and Optimization* **61**(4), 2682–2707 (2023)
23. A. Krause, A. Singh, C. Guestrin, Near-optimal sensor placements in gaussian processes: Theory, efficient algorithms and empirical studies. *Journal of Machine Learning Research* **9**, 235–284 (2008)
24. B. Mu, A.a. Agha-mohammadi, L. Paull, M. Graham, J. How, J. Leonard, Two-stage focused inference for resource-constrained collision-free navigation, in *Robotics: Science and Systems Conference* (Rome, Italy, 2015)
25. E. Paraskevas, D. Maity, J.S. Baras, Distributed energy-aware mobile sensor coverage: A game theoretic approach, in *American Control Conference (ACC)* (Boston, MA, USA, 2016), pp. 6259–6264
26. D. Maity, J.S. Baras, Dynamic, optimal sensor scheduling and value of information, in *18th International Conference on Information Fusion* (Washington, DC, USA, 2015), pp. 239–244
27. M.A. Demetriou, I.I. Hussein, Estimation of spatially distributed processes using mobile spatially distributed sensor network. *SIAM Journal on Control and Optimization* **48**(1), 266–291 (2009)
28. J. Fang, H. Zhang, R.V. Cowlagi, Interactive route-planning and mobile sensing with a team of robotic vehicles in an unknown environment, in *AIAA Scitech 2021 Forum* (2021)
29. F. Bourgault, A. Makarenko, S. Williams, B. Grocholsky, H. Durrant-Whyte, Information based adaptive robotic exploration, in *IEEE/RSJ International Conference on Intelligent Robots and Systems* (Lausanne, Switzerland, 2002)
30. A. Asgharivaskasi, N. Atanasov, Semantic octree mapping and Shannon mutual information computation for robot exploration. *IEEE Transactions on Robotics* **39**(3), 1910–1928 (2023)
31. B. Charrow, S. Liu, V. Kumar, N. Michael, Information-theoretic mapping using Cauchy-Schwarz quadratic mutual information, in *IEEE International Conference on Robotics and Automation (ICRA)* (Seattle, WA, USA, 2015), pp. 4791–4798
32. A. Elfes, Sonar-based real-world mapping and navigation. *IEEE Transactions on Robotics and Automation* **3**(3), 249–265 (1987)
33. H.P. Moravec, Sensor fusion in certainty grids for mobile robots. *AI Magazine* **9**(2), 61–74 (1988)
34. R.V. Cowlagi, P. Tsiotras, Multiresolution motion planning for autonomous agents via wavelet-based cell decompositions. *IEEE Transactions on Systems, Man, and Cybernetics, Part B (Cybernetics)* **42**(5), 1455–1469 (2012)

- 35. G.K. Kraetzschmar, G.P. Gassull, K. Uhl, Probabilistic quadrees for variable-resolution mapping of large environments. *IFAC Proceedings Volumes* **37**(8), 675–680 (2004)
- 36. D.T. Larsson, D. Maity, P. Tsiotras, Information-theoretic abstractions for planning in agents with computational constraints. *IEEE Robotics and Automation Letters* **6**(4), 7651–7658 (2021)
- 37. E.W. Dijkstra, A note on two problems in connexion with graphs. *Numerische Mathematik* **1**(1), 269–271 (1959)