
Near-Optimal Multi-Agent Learning for Safe Coverage Control

Manish Prajapat
ETH Zurich
manishp@ai.ethz.ch

Matteo Turchetta
ETH Zurich
matteotu@inf.ethz.ch

Melanie N. Zeilinger[†]
ETH Zurich
mzeilinger@ethz.ch

Andreas Krause[†]
ETH Zurich
krausea@ethz.ch

Abstract

In multi-agent coverage control problems, agents navigate their environment to reach locations that maximize the coverage of some density. In practice, the density is rarely known *a priori*, further complicating the original NP-hard problem. Moreover, in many applications, agents cannot visit arbitrary locations due to *a priori* unknown safety constraints. In this paper, we aim to efficiently learn the density to approximately solve the coverage problem while preserving the agents' safety. We first propose a conditionally linear submodular coverage function that facilitates theoretical analysis. Utilizing this structure, we develop MACOPT, a novel algorithm that efficiently trades off the exploration-exploitation dilemma due to partial observability, and show that it achieves sublinear regret. Next, we extend results on single-agent safe exploration to our multi-agent setting and propose SAFEMAC for safe coverage and exploration. We analyze SAFEMAC and give first of its kind results: near optimal coverage in finite time while provably guaranteeing safety. We extensively evaluate our algorithms on synthetic and real problems, including a biodiversity monitoring task under safety constraints, where SAFEMAC outperforms competing methods.

1 Introduction

In multi-agent coverage control (MAC) problems, multiple agents coordinate to maximize coverage over some spatially distributed events. Their applications abound, from collaborative mapping [1], environmental monitoring [2], inspection robotics [3] to sensor networks [4]. In addition, the coverage formulation can address core challenges in cooperative multi-agent RL [5, 6], e.g., *exploration* [7], by providing high-level goals. In these applications, agents often encounter safety constraints that may lead to critical accidents when ignored, e.g., obstacles [8] or extreme weather conditions [9, 10].

Deploying coverage control solutions in the real world presents many challenges: (i) for a given density of relevant events, this is an *NP hard problem* [11]; (ii) such *density* is *rarely known* in practice [2] and must be learned from data, which presents a complex active learning problem as the quantity we measure (the density) differs from the one we want to optimize (its coverage); (iii) agents often operate under *safety-critical* conditions, [8–10], that may be *unknown a priori*. This requires cautious exploration of the environment to prevent catastrophic outcomes. While prior work addresses subsets of these challenges (see Section 7), we are not aware of methods that address them jointly.

[†] Joint supervision. Code available at <https://github.com/manish-pra/SafeMaC>

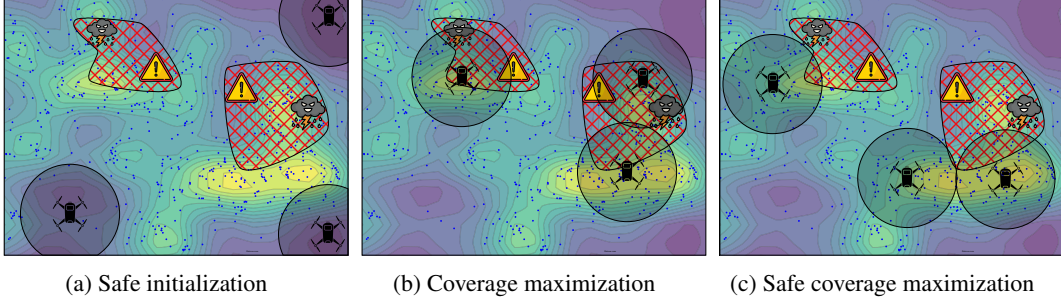


Figure 1: The three drones aim to maximize the gorilla nests’ coverage (grey shaded circle) while avoiding unsafe extreme weather zones (red cross pattern). The contours (yellow is high and purple is low) represent the density of gorilla nests (blue dots). The density and the constraint are a-prior unknown. a) To be safe, drones apply a conservative strategy and do not explore, which results in poor coverage. In b), the drones maximize coverage but get destroyed in extreme weather. c) shows SAFEMAC solution. The drones strike a balance, trading off between learning the density and the constraints, and thus achieve near-optimal coverage while always being safe.

This work makes the following contributions toward efficiently solving safe coverage control with *a-priori* unknown objectives and constraints. **Firstly**, we model this multi-agent learning task as a *conditionally linear* coverage function. We use the *monotonicity* and the *submodularity* of this function to propose MACOPT, a new algorithm for the unconstrained setting that enjoys sublinear cumulative regret and efficiently recommends a near-optimal solution. **Secondly**, we extend GoOSE [12], an algorithm for single agent safe exploration, to the multi-agent case. Combining our extension of GoOSE with MACOPT, we propose SAFEMAC, a novel algorithm for safe multi-agent coverage control. We analyze it and show it attains a near-optimal solution in a finite time. **Finally**, we demonstrate our algorithms on a synthetic and two real world applications: safe biodiversity monitoring and obstacle avoidance. We show SAFEMAC finds better solutions than algorithms that do not actively explore the feasible region and is more sample efficient than competing near-optimal safe algorithms.

2 Problem Statement

We present the safety-constrained multi-agent coverage control problem (Fig. 1) that we aim to solve.

Coverage control. Coverage control models situations where we want deploy a swarm of dynamic agents to maximize the coverage of a quantity of interest, see Fig. 1. Formally, given a finite¹ set of possible locations V , the goal of coverage control is to maximize a function $F: 2^V \rightarrow \mathbb{R}$ that assigns to each subset, $X \subseteq V$, the corresponding coverage value. For N agents, the resulting problem is $\arg \max_{X: |X| \leq N} F(X)$. The discrete domain V can be represented by a graph, where nodes represent locations in the domain, and an edge connects node v to v' if the agent can go from v to v' . This corresponds to a deterministic MDP where locations are states and edges represent transitions.

Sensing region. Depending on the application, we may use different definitions of F . Here, we model cases where agent i at location x^i covers a limited sensing region around it, D^i . While D^i can be any connected subset of V , in practice it is often a ball centered at x^i . Given a function $\rho: V \rightarrow \mathbb{R}$ denoting the density of a quantity of interest at each $v \in V$, our coverage objective is

$$F(X; \rho, V) = \sum_{x^i \in X} \sum_{v \in D^{i-}} \rho(v) / |V|, \quad (1)$$

where $D^{i-} := D^i \setminus D^{1:i-1}$ indicates the elements in V covered by agent i but not agents $1: i-1$, $D^{1:i-1} = \cup_{j=1}^{i-1} D^j$ and $|V|$ denotes cardinality of the domain V .

Safety. In many real-world problems, agents cannot go to arbitrary locations due to safety concerns. To model this, we introduce a constraint function $q: V \rightarrow \mathbb{R}$ and we consider safe all locations v satisfying $q(v) \geq 0$. Such constraint restricts the space of possible solutions of our problem in two ways. First, it prevents agents from monitoring from unsafe locations. Second, depending on its dynamics, agent i may be unable to safely reach a disconnected safe area starting from x_0^i , see

¹Continuous domains can be handled via discretization

Appendix A.3. We denote with $\bar{R}_{\epsilon_q}(\{x_0^i\})$ the largest safely reachable region starting from x_0^i and with B a collection of batches of agents such that all agents in the same batch B share the same safely reachable set, $\forall i, j \in B: \bar{R}_{\epsilon_q}(\{x_0^i\}) \cap \bar{R}_{\epsilon_q}(\{x_0^j\}) \neq \emptyset$, see Appendix A for formal definitions. Based on this, we define the safely reachable control problem

$$\sum_{B \in \mathcal{B}} \max_{X^B \in \bar{R}_{\epsilon_q}(X_0^B)} F(X^B; \rho, \bar{R}_{\epsilon_q}(X_0^B)), \quad (2)$$

where $X_0^B = \{x_0^i\}_{i \in B}$ are the starting locations of all agents in B and $\bar{R}_{\epsilon_q}(X_0^B) = \bigcup_{i \in B} \bar{R}_{\epsilon_q}(\{x_0^i\})$ indicates the largest safely reachable region from any point x_0^i for all i in B (since the agents have the same dynamics, $\bar{R}_{\epsilon_q}(X_0^B) = \bar{R}_{\epsilon_q}(\{x_0^i\})$, $\forall i \in B$). In safety-critical monitoring, there may be unreachable safe regions. However, since agents should be able to collect measurements if required, we focus only on covering the safely reachable region.

Unknown density and constraint. In practice, the density ρ and the constraint q are often unknown *a priori*. However, the agents can iteratively obtain noisy measurements of their values at target locations. We consider synchronous measurements, i.e., we wait until all agents have collected the desired measurement for the current iteration before moving to the next one. Here, we focus on the high-level problem of choosing informative locations, rather than the design of low-level motion planning². Therefore, our goal is to find an approximate solution to the problem in Eq. (2) preserving safety throughout exploration, i.e., at every location visited by the agents, while taking as few measurements as possible in case the dynamics of the agents are deterministic and known as in [12].

3 Background

This section presents foundational ideas that our method builds on. In particular, it discusses (i) monotone submodular functions and (ii) previous work on single-agent safe exploration.

Submodularity. Optimizing a function defined over the power set of a finite domain, V , scales combinatorially with the size of V in general. In special cases, we can exploit the structure of the objective to find approximate solutions efficiently. Monotone submodular functions are one example of this.

A set function $F : 2^V \rightarrow \mathbb{R}$ is *monotone* if for all $A \subseteq B \subset V$ we have $F(A) \leq F(B)$. It is *submodular* if $\forall A \subseteq B \subseteq V, v \in V \setminus B$, we have, $F(A \cup \{v\}) - F(A) \geq F(B \cup \{v\}) - F(B)$. In coverage control, this means adding v to A yields at least as much increase in coverage than adding v to B , if $A \subseteq B$. Crucially, [13] guarantees that the greedy algorithm produces a solution within a factor of $(1 - 1/e)$ of the optimal solution for problems of the type $\arg \max_{X: |X| \leq N} F(X; \rho, V)$, when F is monotone and submodular. In practice, the greedy algorithm often outperforms this worst-case guarantee [14] and guaranteeing a solution better than $(1 - 1/e)$ factor is NP hard [15].

The coverage function in Eq. (1) is a conditionally linear, monotone and submodular function (proof in Appendix B), which lets us use the results above to design our algorithm for safe coverage control.

Goal-oriented safe exploration. GoOSE [12] is a single-agent safe exploration algorithm that extends unconstrained methods to safety-critical cases. Concretely, it maintains under- and over-approximations of the feasible set, called pessimistic and optimistic safe sets. It preserves safety by restricting the agent to the pessimistic safe set. It efficiently explores the objective by letting the original unconstrained algorithm recommend locations within the optimistic safe set. If such recommendations are provably safe, the agent evaluates the objective there. Otherwise, it evaluates the constraint at a sequence of safe locations to prove that such recommendation is either safe, which allows it to evaluate the objective, or unsafe, which triggers the unconstrained algorithm to provide a new recommendation.

Assumptions. To guarantee safety, GoOSE makes two main assumptions. First, it assumes there is an initial set of safe locations, X_0 , from where the agent can start exploring. Second, it assumes the constraint is sufficiently well-behaved, so that we can use data to infer the safety of unvisited locations. Formally, it assumes the domain V is endowed with a positive definite kernel $k^q(\cdot, \cdot)$, and that the constraint's norm in the associated *Reproducing Kernel Hilbert Space* [16] is bounded, $\|q\|_{k^q} \leq B_q$. This lets us use Gaussian Processes (GPs) [17] to construct high-probability confidence intervals for q . We specify the GP prior over q through a mean function, which we assume to be

²Agents can use their transition graph to find a path between two goals. In a continuous domain, the path can be tracked with a controller (e.g., MPC)

zero everywhere w.l.o.g., $\mu(v) = 0, \forall v \in V$, and a kernel function, k , that captures the covariance between different locations. If we have access to T measurements, at $V_T = \{v_t\}_{t=1}^T$ perturbed by i.i.d. Gaussian noise, $y_T = \{q(v_t) + \eta_t\}_{t=1}^T$ with $\eta_t \sim \mathcal{N}(0, \sigma^2)$, we can compute the posterior mean and covariance over the constraint at unseen locations v, v' as $\mu_T(v) = k_T^\top(v)(K_T + \sigma^2 I)^{-1} y_T$ and $k_t(v, v') = k(v, v') - k_T^\top(v)(K_T + \sigma^2 I)^{-1} k_T(v')$, where $k_T(v) = [k(v_1, v), \dots, k(v_T, v)]^\top$, K_T is the positive definite kernel matrix $[k(v, v')]_{v, v' \in V_T}$ and $I \in \mathbb{R}^{T \times T}$ denotes the identity matrix.

In this work, we make the same assumptions about the safe seed and the regularity of q and ρ .

Approximations of the feasible set. Based on the GP posterior above, GoOSE builds monotonic confidence intervals for the constraint at each iteration t as $l_t^q(v) := \max\{l_{t-1}^q(v), \mu_{t-1}^q(v) - \beta_t^q \sigma_{t-1}^q(v)\}$ and $u_t^q(v) := \min\{u_{t-1}^q(v), \mu_{t-1}^q(v) + \beta_t^q \sigma_{t-1}^q(v)\}$, which contain the true constraint function for every $v \in V$ and $t \geq 1$, with high probability if β_t^q is selected as in [18] or Section 5. GoOSE uses these confidence intervals within a set $S \subseteq V$ together with the L_q -Lipschitz continuity of q , to define operators that determine which locations are safe in plausible worst- and best-case scenarios,

$$p_t(S) = \{v \in V, \exists z \in S : l_t^q(z) - L_q d(v, z) \geq 0\}, \quad (3)$$

$$o_t^{\epsilon_q}(S) = \{v \in V, \exists z \in S : u_t^q(z) - \epsilon_q - L_q d(v, z) \geq 0\}. \quad (4)$$

Notice that the pessimistic operator relies on the lower bound, l^q , while the optimistic one on the upper bound, u^q . Moreover, the optimistic one uses a margin ϵ_q to exclude "barely" safe locations as the agent might get stuck learning about them. Finally, to disregard locations the agent could not safely reach or from where it could not safely return, GoOSE introduces the $R^{\text{ergodic}}(\cdot, \cdot)$ operator. $R^{\text{ergodic}}(p_t(S), S)$ indicates locations in S or locations in $p_t(S)$ reachable from S and from where the agent can return to S along a path contained in $p_t(S)$. Combining $p_t(S)$ and $R^{\text{ergodic}}(\cdot, \cdot)$, GoOSE defines the pessimistic and ergodic operator $\hat{P}_t(\cdot)$, which it uses to update the pessimistic safe set. Similarly, it defines $\hat{O}_t(\cdot)$ using $o_t^{\epsilon_q}(\cdot)$ to compute the optimistic safe set.

4 MACOPT and SAFEMAC

This section presents MACOPT and SAFEMAC, our algorithms for unconstrained and safety-constrained multi-agent coverage control, which we then formally analyze in Section 5.

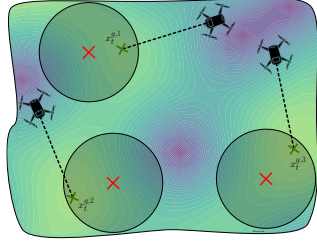
4.1 MACOPT: unconstrained multi-agent coverage control

Greedy sensing regions. In sequential optimization, it is crucial to balance exploration and exploitation. GP-UCB [19] is a theoretically sound strategy to strike such a trade-off that works well in practice. Agents evaluate the objective at locations that maximize an upper confidence bound over the objective given by the GP model such that locations with either a high posterior mean (exploitation) or standard deviation (exploration) are visited. We construct a valid upper confidence bound for the coverage $F(X)$ starting from our confidence intervals on ρ , by replacing the true density ρ with its upper bound u_t^ρ in Eq. (1). Next, we apply the greedy algorithm to this upper bound (Line 3 of Algorithm 1) to select N candidate locations for evaluating the density. However, this simple exploration strategy may perform poorly, due to the fact that in order to reduce the uncertainty over the coverage F at X , we must learn the density ρ at all locations inside the sensing region, $\bigcup_{i=1}^N D^i$, rather than simply at X . It is a form of partial monitoring [20], where the objective F differs from the quantity we measure, i.e., the density ρ . Next, we explain how to choose locations where to observe the density for a given X .

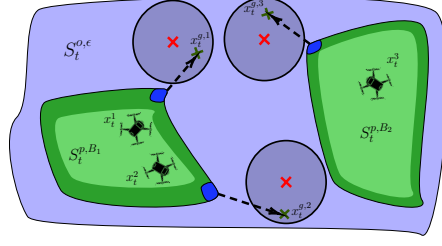
Uncertainty sampling. Given location assignments X for the agents, we measure the density to efficiently learn the function $F(X)$. Intuitively, agent i observes the density where it's most uncertain within the area it covers that is not covered by agents $\{1, \dots, i-1\}$, i.e., D_t^{i-} (Line 4 of Alg. 2, Fig. 2a).

Stopping criterion. The algorithm terminates when a near-optimal solution is achieved. Intuitively, this occurs when the uncertainty about the coverage value of the greedy recommendation is low. Formally, we require the sum of the uncertainties over the sampling targets to be below a threshold, i.e., $w_t = \sum_{i=1}^N u_{t-1}^\rho(x_t^{g,i}) - l_{t-1}^\rho(x_t^{g,i}) \leq \epsilon_\rho$ (Line 3 of Algorithm 2). Importantly, this stopping criterion requires the confidence intervals to shrink only at regions that potentially maximize the coverage.

MACOPT. Now, we introduce MACOPT in Algorithm 2. At round t , we select the sensing locations for the agents, X_t , by greedily optimizing the upper confidence bound of the coverage. Then, each



(a) Uncertainty sampling in MACOPT



(b) Illustration of multi-agent GoOSE

Figure 2: a) The contours represent the density uncertainty, and the red \times 's correspond to the maximum coverage locations evaluated by the GREEDY Algorithm 1. While these locations maximize coverage, they may not be informative about the coverage since the uncertainty can be low. Therefore, the agents collect measurements at the maximum uncertainty of the density in a disc (green \times 's, $x_t^{g,i}$), also known as uncertainty sampling. b) In a constrained environment, SAFEMAC evaluates $x_t^{g,i}$ for all agents in the optimistic set S_t^{o, ϵ_q} (violet) and set it as a next goal. It forms an expander region (dark blue) to safely expand the pessimistic safe set S_t^p (green) toward the goal.

agent i collects noisy density measurements at the points of highest uncertainty within D_t^{i-} . Finally, we update our GP over the density and, if the sum of maximum uncertainties within each sensing region is small, we stop the algorithm.

4.2 SAFEMAC: safety-constrained multi-agent coverage control

Intuition. We adopt a perspective similar to GoOSE as we separate the exploration of the safe set from the maximization of the coverage. Given an over and under approximation of the safe set (whose computation is discussed later), we want to explore optimistically optimal goals for each agent, similar to MACOPT. To this end, we find the maximizers of the density upper bound in the optimistic safe set with the GREEDY algorithm. Then, we define sampling goals to learn the coverage at those locations.

Phases of SAFEMAC. Coverage values depend both on the density and the feasible region (Eq. (2)). Thus, there are two sensible sampling goals given a disk assignment: i) *optimistic coverage*: if we are uncertain about the density within the disks, we target locations with the highest density uncertainty (Line 6 of Algorithm 4); ii) *optimistic exploration*: if we know the density within the disk but there are locations under it that we cannot classify as either safe (in S^p) or unsafe (in $V \setminus S^{o, \epsilon_q}$), we target those with the highest constraint uncertainty among them (Line 8). If all the goal locations are safe with high probability, which can only happen during *optimistic coverage*, we safely evaluate the density there (Line 19). Otherwise, we explore the constraint with a goal directed strategy that aims at classifying them as either safe or unsafe similar to GoOSE (Line 9-12). In case this changes the topological connection of the optimistic feasible set, we recompute the disks as this may change GREEDY's output (Line 15-17). We repeat this loop until we know the feasibility of all the points under the disks recommended by GREEDY and their density uncertainty is low (Line 4). Next, we explain how the multiple agents coordinate their individual safe regions to evaluate a goal (MACOPT *in batches*), how the agents progress toward their goals (*safe expansion*) and finally we describe SAFEMAC *convergence*.

MACOPT in batches. In the multi-agent setting of GoOSE (see Fig. 2b), each agent i maintains $S_t^{p,i}$ a pessimistic (or $S_t^{o, \epsilon_q, i}$ an optimistic) belief of the safe locations, obtained by iteratively applying $\tilde{P}_t(\cdot)$ the pessimistic (or $\tilde{O}_t(\cdot)$ the optimistic) ergodic operators (see Section 3) to the previous pessimistic belief $S_{t-1}^{p,i}$ (Line 11 of Algorithm 4). Since the agents cannot navigate to an arbitrary location in the constrained case, SAFEMAC computes coverage maximizers on a restricted region, obtained by ignoring the known unsafe locations. To denote such a restricted region, we define a union set $S_t^{u,i} := S_t^{o, \epsilon_q, i} \cup S_t^{p,i}$, which is the largest set known to be optimistically or pessimistically safe up to time t . Moreover, if the agents are topologically disconnected, they cannot travel from one safe region to another and the best strategy for any batch of agents is to maximize coverage locally. For this, we form a collection of batches \mathcal{B}_t , such that any batch $B \in \mathcal{B}_t$ contains agents that lie in topologically connected regions determined by the union set (Line 13-14). SAFEMAC computes a GREEDY solution for each $B \in \mathcal{B}_t$ in their corresponding $S_t^{u,B} := \cup_{i \in B} S_t^{u,i}$. This is the largest set where the agents can find an

Algorithm 1 Greedy UCB (GREEDY)

```

1: Inputs  $u_{t-1}^\rho, l_{t-1}^\rho, B, S_t^u$ 
2: for  $i = 1, 2, \dots, |B|$  do
3:    $x_t^i \leftarrow \arg \max_{x^i} \sum_{v \in D^i \setminus D_{t-1}^{1:i-1} \cap S_t^u} u_{t-1}^\rho(v)$ 
4:    $x_t^{g,i} \leftarrow \arg \max_{v \in D^i \setminus D_{t-1}^{1:i-1} \cap S_t^u} u_{t-1}^\rho(v) - l_{t-1}^\rho(v)$ 
5:  $w_t \leftarrow \sum_{i=1}^{|B|} u_{t-1}^\rho(x_t^{g,i}) - l_{t-1}^\rho(x_t^{g,i})$ 
6: Return  $X_t^B, w_t$ 

```

Algorithm 2 MACOPT

```

1: Inputs  $X_0, \epsilon_\rho, V, GP_\rho, t \leftarrow 1$ 
2:  $X_1, w_1 \leftarrow \text{GREEDY}(u_0^\rho, l_0^\rho, [N], V)$ 
3: while  $w_t > \epsilon_\rho$  do
4:    $\forall i, x_t^{g,i} \leftarrow \arg \max_{v \in D_t^{i-}} u_{t-1}^\rho(v) - l_{t-1}^\rho(v)$ 
5:    $\forall i, y_{\rho_t}^i = \rho(x_t^{g,i}) + \eta_\rho$ , Update GP
6:    $t \leftarrow t + 1$ 
7:    $X_t, w_t \leftarrow \text{GREEDY}(u_{t-1}^\rho, l_{t-1}^\rho, [N], V)$ 
8: Recommend  $X_t$ 

```

Algorithm 3 Safe Expansion (SE)

```

1: Inputs  $S_t^{o,\epsilon_q}, S_t^p, x_t^g$ 
2:  $A_t(p) \leftarrow \{v \in S_t^{b,\epsilon_q} \mid p_t(S_t^p) \mid h(v) = p\}$ 
3:  $W_t^{\epsilon_q} \leftarrow \{v \in S_t^p \mid u_t^q(v) - l_t^q(v) > \epsilon_q\}$ 
4:  $\alpha^* \leftarrow \max \alpha \text{ s.t. } |G_t^{\epsilon_q}(\alpha)| > 0$ 
5: if Optimization problem feasible then
6:    $v_t \leftarrow \arg \max_{v \in G_t^{\epsilon_q}(\alpha^*)} u_t^q(v) - l_t^q(v)$ 
7:   Update GP with  $y_t = q(v_t) + \eta_q$ 

```

Algorithm 4 SAFEMAC

```

1: Inputs  $X_0, L_q, \epsilon_\rho, V, GP_\rho, GP_q$ 
2:  $\forall i, S_0^{p,i} \leftarrow X_0, S_0^{o,\epsilon_q,i} \leftarrow V, t \leftarrow 1$ 
3:  $X_1, w_1 \leftarrow \text{GREEDY}(u_0^\rho, l_0^\rho, [N], V)$ 
4: while  $\forall i, (S_{t-1}^{o,\epsilon_q,i} \setminus S_{t-1}^{p,i}) \cap D_t^i \neq \emptyset$  or  $w_t > \epsilon_\rho$  do
5:   if  $w_t > \epsilon_\rho$  then
6:      $\forall i, x_t^{g,i} \leftarrow \arg \max_{v \in D_t^{i-}} u_{t-1}^\rho(v) - l_{t-1}^\rho(v)$ 
7:   else
8:      $\forall i, x_t^{g,i} \leftarrow \arg \max_{v \in (S_{t-1}^{o,\epsilon_q,i} \setminus S_{t-1}^{p,i}) \cap D_t^i} u_{t-1}^q(v) - l_{t-1}^q(v)$ 
9:   if  $\exists i \in [N], x_t^{g,i} \notin S_t^{p,i}$  then
10:     $\text{SE}(S_{t-1}^{o,\epsilon_q,i}, S_{t-1}^{p,i}, x_t^{g,i}, \forall i : x_t^{g,i} \notin S_t^{p,i})$ 
11:     $S_t^{p,i} \leftarrow \tilde{P}_t(S_{t-1}^{p,i}), S_t^{o,\epsilon_q,i} \leftarrow \tilde{O}_t^{\epsilon_q}(S_{t-1}^{p,i}), \forall i$ 
12:     $t \leftarrow t + 1$ 
13:     $\forall i, \mathcal{B}'_t(i) = \{j \in [N] \mid S_t^{u,i} \cap S_t^{u,j} \neq \emptyset\}$ 
14:     $\mathcal{B}_t = \bigcup_{i \in [N]} \mathcal{B}'_t(i)$ 
15:    if for any  $B \in \mathcal{B}_t, S_t^{u,B} \neq S_{t-1}^{u,B}$  then
16:       $X_t, w_t \leftarrow \text{GREEDY}(u_{t-1}^\rho, l_{t-1}^\rho, B, S_t^{u,B})$ 
17:       $\forall i, x_t^{g,i} \leftarrow \arg \max_{v \in D_t^{i-}} u_{t-1}^\rho(v) - l_{t-1}^\rho(v)$ 
18:    if  $\forall i, x_t^{g,i} \in S_t^{p,i}$  and  $w_t > \epsilon_\rho$  then
19:       $\forall i, y_{\rho_t}^i = \rho(x_t^{g,i}) + \eta_\rho$ 
20:      Update GP i.e., compute  $u_t^\rho, l_t^\rho$ 
21:       $t \leftarrow t + 1$ 
22:       $X_t, w_t \leftarrow \text{GREEDY}(u_{t-1}^\rho, l_{t-1}^\rho, B, S_t^{u,B})$ 
23: Recommend  $X_t$ 

```

optimistically safe path to travel. Analogous to \mathcal{B}_t , we define \mathcal{B}_t^p as collection of batches where any $B \in \mathcal{B}_t^p$ contains agents which are topologically connected in pessimistic set and $S_t^{p,B} := \cup_{i \in B} S_t^{p,i}$.

Safe expansion. Safe expansion is the sub-routine inspired by GoOSE for goal-oriented exploration of the safe set that we use to learn about the feasibility of sampling targets. It uses a heuristic h to assign priority scores p to points that are optimistically but not pessimistically safe. Those determine locations whose feasibility is relevant to learn that of the sampling targets (Line 2 of Algorithm 3). A simple and effective choice for the heuristic is the inverse of the distance to the targets. Then, it identifies safe locations where the constraint is not yet known ϵ_q -accurately (Line 3). Among them, it determines the α -immediate expanders, i.e., those that could potentially add locations with priority α to the pessimistic set, $G_t^{\epsilon_q}(\alpha) = \{v \in W_t^{\epsilon_q} \mid \exists z \in A_t(\alpha) : u_t^q(v) - L_q d(v, z) \geq 0\}$. In Line 4, it selects the non-empty α -expander set with the highest priority. In Line 6 - 7, the agent evaluates the constraint at the location with the highest uncertainty in this set (see [12] for details).

SAFEMAC convergence. The *optimistic coverage* phase switches to *optimistic exploration* phase, when density uncertainty under the disks is low ($w_t \leq \epsilon_\rho$). In the exploration, either the topological connection of the optimistic feasible set changes or will classify the uncertain region as pessimistically safe. In the former case, SAFEMAC will recompute a new coverage location and switch to the coverage phase. Alternatively, if the uncertain region is pessimistically safe, SAFEMAC has converged since the density uncertainty in the exploration phase is already low. The phases show an interesting dynamics; SAFEMAC continuously iterates between the *optimistic exploration* and the *optimistic coverage* phase until we know about the feasibility of the disk and their uncertainty is low. In the worst case, SAFEMAC might explore the entire environment. In this case the sample complexity will be similar to a two-stage algorithm, where we explore the whole domain and then optimize coverage in the resulting known environment. However, in practice, SAFEMAC is much better than this worst case.

5 Analysis

We now analyze MACOPT's convergence and SAFEMAC's optimality and safety properties.

MACOPT. To measure the progress of MACOPT, we study its regret, i.e., the difference between its solution and the one we could find if we knew the true density. Since control coverage consists in maximizing a monotone submodular function, we cannot efficiently compute the true optimum even for known densities. However, we can efficiently find a solution that is at least $(1 - 1/e)$ within the optimum. Thus, we quantify performance using the following notion of cumulative regret,

$$Reg_{act}(T) = \left(1 - \frac{1}{e}\right) \sum_{t=1}^T F(X_\star; \rho, V) - \sum_{t=1}^T F(X_t; \rho, V), \quad (5)$$

where $F(X_\star; \rho, V)$ is the optimal coverage. We now state one of our main results, which guarantees that the cumulative regret of MACOPT grows sublinearly in time (proof in Appendix D).

Theorem 1. *Let $\delta \in (0, 1)$, $\beta_t^{\rho^{1/2}} = B_\rho + 4\sigma_\rho \sqrt{\gamma_{Nt}^\rho + \ln(1/\delta)}$ and $C_D = \max_{x^i \in V} |D^i|/|V| \leq 1$. With probability at least $1 - \delta$, MACOPT's regret defined in Eq. (5) is bounded by $\mathcal{O}(\sqrt{T \beta_T^\rho \gamma_{NT}^\rho})$,*

$$Pr \left\{ Reg_{act}(T) \leq \sqrt{\frac{8C_D N T \beta_T^\rho \gamma_{NT}^\rho}{\log(1 + N \sigma_\rho^{-2})}} \right\} \geq 1 - \delta. \quad (6)$$

The proof of 1 builds on two key ideas. First, we exploit the conditional linearity of the submodular objective to bound the cumulative regret defined in Eq. (5) with a sum of per agent regrets. Secondly, we bound the per agent regret with the information capacity γ_{NT}^ρ , a quantity that measures the largest reduction in uncertainty about the density that can be obtained from NT noisy evaluations of it. Since γ_{NT}^ρ [21] grows sublinearly with T for commonly used kernels, so does MACOPT's regret in Eq. (6). The immediate corollary of the above theorem, when the MACOPT stopping criteria is reached (Line 3 of Algorithm 2) guarantees a near optimal solution up to ϵ_ρ precision.

Corollary 1. *Let t_ρ^\star be the smallest integer, such that $\frac{t_\rho^\star}{\beta_{t_\rho^\star} \gamma_{Nt_\rho^\star}^\rho} \geq \frac{8C_D^2 N^2}{\log(1 + N \sigma_\rho^{-2}) \epsilon_\rho^2}$, then there exists a $t < t_\rho^\star$ such that w.h.p, MACOPT terminates and achieves, $F(X_t; \rho, V) \geq (1 - \frac{1}{e})F(X_\star; \rho, V) - \epsilon_\rho$.*

SAFEMAC. This section presents our main result for safety-constrained multi-agent coverage control. In particular, Theorem 2 (proof in Appendix E) guarantees that SAFEMAC safely achieves near-optimal safe coverage in finite time.

Theorem 2. *Let $\delta \in (0, 1)$, $\epsilon_\rho \geq 0$, $\|\rho\|_{k^\rho} \leq B_\rho$, $\beta_t^{\rho^{1/2}} = B_\rho + 4\sigma_\rho \sqrt{\gamma_{Nt}^\rho + 1 + \ln(1/\delta)}$, γ_{Nt}^ρ denote the information capacity associated with the kernel k^ρ . Let $q(\cdot)$ be L_q -Lipschitz continuous and $\epsilon_q, \beta_t^q, \gamma_{Nt}^q$ be defined analogously. Given $X_0 \neq \emptyset$, $q(x_0^i) \geq 0$ for all $i \in [N]$. Then, for any heuristic $h_t : V \rightarrow \mathbb{R}$, with probability at least $1 - \delta$, we have $q(x) \geq 0$, for any x along the state trajectory pursued by any agent in SAFEMAC. Moreover, let t_ρ^\star be the smallest integer such that $\frac{t_\rho^\star}{\beta_{t_\rho^\star} \gamma_{Nt_\rho^\star}^\rho} \geq \frac{8C_D^2 N^2}{\log(1 + N \sigma_\rho^{-2}) \epsilon_\rho^2}$, with $C_D = \max_{x^i \in V} \frac{|D^i|}{|V|} \leq 1$ and let t_q^\star be the smallest integer such that $\frac{t_q^\star}{\beta_{t_q^\star} \gamma_{Nt_q^\star}^q} \geq \frac{C |\bar{R}_0(X_0)|}{\epsilon_q^2}$, with $C = 8/\log(1 + \sigma_q^{-2})$ then, there exists $t \leq t_q^\star + t_\rho^\star$, such that with probability at least $1 - \delta$,*

$$\sum_{B \in \mathcal{B}_t} F(X_t^B; \rho, \bar{R}_0(X_0^B)) \geq \left(1 - \frac{1}{e}\right) \sum_{B \in \mathcal{B}} F(X_\star^B; \rho, \bar{R}_{\epsilon_q}(X_0^B)) - \epsilon_\rho. \quad (7)$$

The theoretical analysis has two components: (i) we show SAFEMAC's coverage is near-optimal at convergence (Lemma 10), and (ii) we prove it converges in finite time. Since SAFEMAC learns the constraint and the density, we must bound the sample complexity for both to prove (ii). For the constraint, we extend the results for single-agent GoOSE to our multi-agent setting (Appendix F). For the density, we use results from Theorem 1 to show that, within a coverage phase, the cumulative regret is sublinear. Next, we use additivity of the information gain (Lemma 13) between any pair of coverage phases to bound the sample complexity of density for the subsequent coverage phases. Combining these results, we obtain Theorem 2.

Intermediate recommendation. Theorem 2 guarantees that SAFEMAC converges to a safe and near-optimal solution. Can it also make sensible recommendations before the stopping criteria are met?

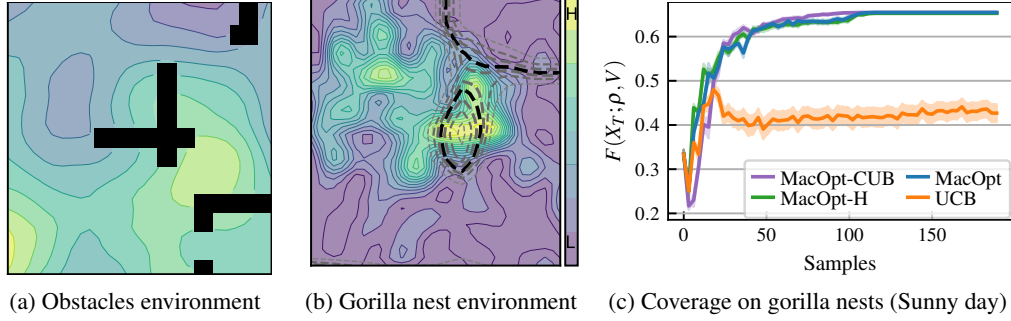


Figure 3: The contours in: a) show the synthetic density and the obstacles marked by the black blocks, b) show the Gorilla nests distribution with weather constraints marked by the black dashed line, and its contours with grey dashed line. c) Compares MACOPT with UCB in the safe gorilla environment. MACOPT does a more principled exploration of the coverage and does not stick to a local minimum.

Ideally, such recommendations should (i) be safely reachable and (ii) ensure a minimum coverage. To satisfy (i), they should be in the pessimistic safe set, S_t^p . To satisfy (ii), their coverage should be computed according to $F(\cdot; l_{t-1}^p, S_t^p)$, i.e., assuming a worst-case density, l_{t-1}^p , and a worst-case feasible set, S_t^p . If the greedy recommendation X_t is in S_t^p , we can recommend it at intermediate steps. However, this is not always the case and we need an alternative. To this end, we compute $X_t^{l,B}$, i.e., the greedy solution w.r.t. the worst-case objective, $F(\cdot; l_{t-1}^p, S_t^{p,B}) \forall B \in \mathcal{B}_t^p$. At any time T , SAFEMAC recommends the best of either strategy up to time T according to the worst-case objective. In Appendix E.1, we show that such recommendation is also near optimal at convergence.

6 Experiments

This section compares MACOPT and SAFEMAC to existing methods (or their extensions) on synthetic and real-world problems. We validate our theoretical claims and observe their superiority. We set $\beta^q = 3$ and $\beta^p = 3$ for all $t \geq 1$, it ensures safety as well as efficient exploration in practice [12]. Experiment details and extended empirical analysis are in Appendix G.

Environments. We perform our experiments with $N = 3$ agents in a 30×30 grid world where states are evenly spaced over $[0, 3]^2$. Each agent’s disk is defined as the region an agent can reach in $r = 5$ steps in the defined grid. We normalize coverage with a maximum value $\sum_{v \in \bar{R}_0(X_0)} \rho(v)/|V|$. Below, we present the 3 environments we consider.

i) In *synthetic data*, both the density ρ and the constrain q are sampled from a GP with zero mean and Matérn Kernel with $\nu = 2.5$, scale $\sigma_k = 1$, and lengthscale $l = 2$. The observations are perturbed by i.i.d noise $\mathcal{N}(0, 10^{-3})$. ii) In *obstacles*, we sample maps with several block-shaped obstacles (Fig. 3a) and we aim to maximize coverage while avoiding dangerous collisions. At v , each agent senses the distance to the nearest obstacle $d_m(v)$, which could be given by sensors such as 1D-Lidars. We use $q'(v) = 1/(1 + \exp(-1.5d_m(v)))$, to map the distance between $[0, 3]$ and saturate the constraint value for large distances, and we set $q(v) = q'(v) - 0.5$ to avoid collisions. The density is sampled from the same GP as the synthetic case. iii) In *gorilla nest*, we simulate a bio-diversity monitoring task, where we aim to cover areas with high density of gorilla nests with a quadrotor in the Kagwene Gorilla Sanctuary (Fig. 3b). Regions affected by adverse weather (e.g. rain and storms) are unsafe for the drone due to higher chances of crashes and should be avoided. As a proxy for bad weather, we use the cloud coverage data over the KGS from OpenWeather [22]. The nest density is obtained by fitting a smooth rate function [23] over Gorilla nest counts [24].

MACOPT. We compare MACOPT to UCB, a baseline that skips the uncertainty sampling step from Section 4.1 and obtains measurements at the centers of the GREEDY sensing regions. We further develop two sample-efficient extensions of MACOPT: i) Correlated upper bound (CUB), a variant of MACOPT that constructs tighter upper confidence bound of the coverage function utilizing the covariance of density, instead of using the sum of density UCB. ii) Hallucinated uncertainty sampling (H), a variant of MACOPT that samples at the most informative location for each agent i , after hallucinating sampling locations of $\{1, \dots, i-1\}$ agents. Please see Appendix D.1 for theoretical analysis. Fig. 3c shows a comparison in the *gorilla* environment on a day of good weather, i.e. when all locations are safe. Here, UCB gets stuck in a local optimum as it does not reduce the uncertainty

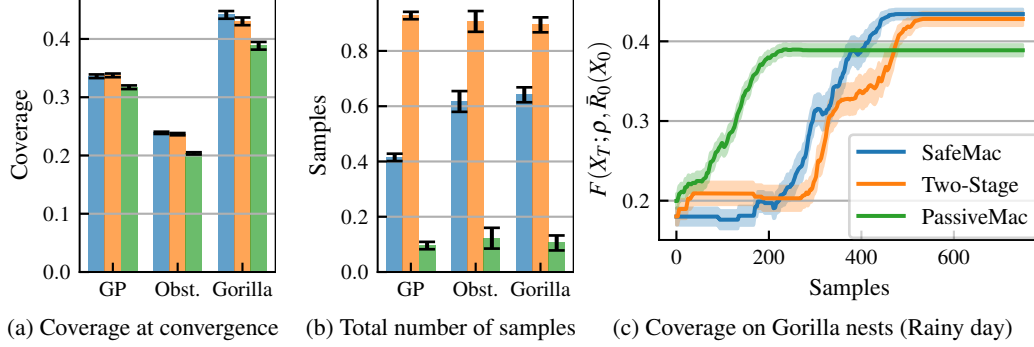


Figure 4: Comparison of SAFEMAC with PASSIVEMAC and Two-Stage in all environments at convergence (a) and (b) and during optimization for the gorilla environment in (c). SAFEMAC trades-off learning about density and constraints, such that it finds a solution comparable to Two-Stage more efficiently, whereas PASSIVEMAC gets stuck in a local optimum.

of the density, whereas MACOPT explores more and achieves a higher coverage value up to 25%. Moreover, variants of MACOPT account for correlation and condition on other agents’ measurement locations, which results in achieving the same coverage but more efficiently.

SAFEMAC. We compare SAFEMAC with two baselines: *i*) a two-stage algorithm [25], that first fully explores the feasible region, and then uses MACOPT to maximize the coverage; *ii*) PASSIVEMAC, a baseline inspired by [26] that runs MACOPT in the pessimistic set and passively measures the constraint in the process. Figs. 4a and 4b show the coverage at convergence and the number of samples to converge for SAFEMAC and the two baselines across all the environments. The results are averaged over 50 instances produced using different seeds and samples for every environment. In Fig. 4b, the y-axis is normalized with the maximum number of samples in the instance and then averaged over all instances. PASSIVEMAC converges quickly but gets stuck in a local optimum as it does not actively explore the constraint. SAFEMAC and Two-Stage converge to much higher coverage values. However, SAFEMAC is up to 50% more sample efficient thanks to its goal-oriented exploration. Fig. 4c shows the coverage value of the intermediate safe recommendations (Section 5) in the *gorilla* environment as a function of the number of samples. It confirms the previous results: SAFEMAC finds solutions comparable to Two-Stage more efficiently and PASSIVEMAC gets stuck in a local optimum.

Scalability. SAFEMAC utilizes the GREEDY algorithm, which is linear in the number of nodes (domain size). In each iteration, SAFEMAC computes a greedy solution N times (one for each agent), which makes it linear in the number of agents. We model density and constraint using GP, which scales cubically with the number of samples. To demonstrate scalability in practice, we conducted experiments with $N = 3, 6, 10, 15$ agents each with domain length of 30, 40, 50 and 60 in Appendix G.1

7 Related work

Our work relates to multiple fields. We highlight the most relevant connections, referencing surveys where possible; an exhaustive overview is beyond the scope of this paper.

Bayesian optimization. In BO, an agent sequentially evaluates a noisy objective, seeking to maximize it [27]. In contrast, the quantity we measure *differs* from our objective. Partial monitoring [28] addresses such issues in an abstract setting [20, 29]. We exploit special structure in our problem. In coverage control with unknown density, this challenge is often addressed by learning the density uniformly over the domain [30, 31]. In contrast, MACOPT learns the density only at promising locations.

Coverage control. MAC with known densities is a well-studied NP hard [32] problem. Many algorithms use efficient heuristics to converge quickly to a local optimum. One popular strategy is Lloyd’s algorithm [33], which has been studied in different settings, e.g., with known densities [34, 35], *a-priori* unknown densities [31, 36–38], using graph neural networks [39], taking into account agent’s dynamics and constraints [40], or in case of non-identical robots [41]. These methods apply to continuous state and action spaces and show convergence to local optima, but lack optimality guarantees [30, 31, 40] and sample complexity bounds. Moreover, their extension to non-convex, disconnected domains is not trivial [42]. Coverage control is also studied in the episodic setting to learn the unknown policy or the environment using deep RL methods [43, 44].

Submodular optimization. Submodular functions are ubiquitous in machine learning [45] as they can be efficiently approximately maximized under different kinds of constraints [46]. For example, the GREEDY algorithm can be used in case of cardinality constraints [13] to maximize quantities like mutual information [47] or weighted coverage functions [15]. Online submodular maximization aims at optimizing unknown submodular functions from noisy measurements [48]. It has multiple applications, including optimization of numerical solvers [49], information gathering [50] and crowd-sourced image collection summarization [51]. Particularly related to ours is the work in [52], which proposes an algorithm for contextual news recommendation for linear user preferences with strong regret guarantees. In contrast to that setting, we consider dynamic agents, safety constraints and partial feedback.

Safety. Depending on the safety formulation and the assumptions, many algorithms have been proposed for safe learning in dynamical systems, e.g., based on model predictive control [53], curriculum learning [54], Lyapunov functions [55, 56], reachability [57], CMDPs [58], behavioral system theory [59], and more [60–62]. Here, we focus on the setting that is most closely related to ours, i.e., one with unknown but sufficiently regular instantaneous constraints that must be satisfied at all times. For stateless problems, e.g. BO, [26, 63] propose algorithms with safety and optimality guarantees with different exploration strategies. For stateful problems, [64] studies the pure exploration case, while [25] extends the two-stage approach from [63]. These approaches may be sample inefficient as they may explore the constraint in regions irrelevant for the objective. GOOSE [12] addresses this problem for both the stateful and stateless setting. The only work in this context that addresses multi-agent problems is [65]. However, their objective differs from ours, and they do not establish safety guarantees.

8 Conclusion

We present two novel algorithms for multi-agent coverage control in unconstrained (MACOPT) and safety critical environments (SAFEMAC). We show MACOPT achieves sublinear cumulative regret, despite the challenge of partial observability. Moreover, we prove SAFEMAC achieves near optimal coverage in finite time while navigating safely. We demonstrate the superiority of our algorithms in terms of sample efficiency and coverage in real-world applications such as safe biodiversity monitoring.

Currently, our algorithms choose informative targets but do not plan informative trajectories, which is crucial in robotics. We aim to address this in future work. Finally, while in many real-world applications the density and the constraints are as regular as assumed here, in some they are not. In these cases, our optimality and safety guarantees would not apply.

Acknowledgements

Manish Prajapat is supported by an ETH AI Center doctoral fellowship. Matteo Turchetta is supported by the Swiss National Science Foundation under NCCR Automation, grant agreement 51NF40 180545. We would like to thank Pawel Czyn for insightful discussions.

References

- [1] Miquel Kegeleirs, Giorgio Grisetti, and Mauro Birattari. Swarm slam: Challenges and perspectives. *Frontiers in Robotics and AI*, 8:618268, 2021.
- [2] Alan Mainwaring, David Culler, Joseph Polastre, Robert Szewczyk, and John Anderson. Wireless sensor networks for habitat monitoring. In *Proceedings of the 1st ACM international workshop on Wireless sensor networks and applications*, pages 88–97, 2002.
- [3] Mahmoud Tavakoli, Gonçalo Cabrita, Ricardo Faria, Lino Marques, and Anibal T de Almeida. Cooperative multi-agent mapping of three-dimensional structures for pipeline inspection applications. *The International Journal of Robotics Research*, 31(12):1489–1503, 2012.
- [4] Bang Wang. *Coverage control in sensor networks*. Springer Science & Business Media, 2010.
- [5] Ryan Lowe, Yi I Wu, Aviv Tamar, Jean Harb, OpenAI Pieter Abbeel, and Igor Mordatch. Multi-agent actor-critic for mixed cooperative-competitive environments. *Advances in neural information processing systems*, 30, 2017.
- [6] Manish Prajapat, Kamyar Azizzadenesheli, Alexander Liniger, Yisong Yue, and Anima Anandkumar. Competitive policy optimization. In *Uncertainty in Artificial Intelligence*, pages 64–74. PMLR, 2021.
- [7] Iou-Jen Liu, Unnat Jain, Raymond A Yeh, and Alexander Schwing. Cooperative exploration for multi-agent deep reinforcement learning. In *International Conference on Machine Learning*, pages 6826–6836. PMLR, 2021.
- [8] Pericle Salvini, Diego Paez-Granados, and Aude Billard. Safety concerns emerging from robots navigating in crowded pedestrian areas. *International Journal of Social Robotics*, 14(2): 441–462, 2022.
- [9] Yuliya Averyanova and E. Znakovskaja. Weather hazards analysis for small uass durability enhancement. In *2021 IEEE 6th International Conference on Actual Problems of Unmanned Aerial Vehicles Development (APUAVD)*, pages 41–44, 2021. doi: 10.1109/APUAVD53804.2021.9615440.
- [10] Mozhou Gao, Chris H Hugenholtz, Thomas A Fox, Maja Kucharczyk, Thomas E Barchyn, and Paul R Nesbit. Weather constraints on global drone flyability. *Scientific reports*, 11(1):1–13, 2021.
- [11] Andreas Krause and Carlos Guestrin. Submodularity and its applications in optimized information gathering. *ACM Trans. Intell. Syst. Technol.*, 2(4), jul 2011. ISSN 2157-6904. doi: 10.1145/1989734.1989736. URL <https://doi.org/10.1145/1989734.1989736>.
- [12] Matteo Turchetta, Felix Berkenkamp, and Andreas Krause. Safe exploration for interactive machine learning. *Advances in Neural Information Processing Systems*, 32, 2019.
- [13] George Nemhauser, Laurence Wolsey, and M. Fisher. An analysis of approximations for maximizing submodular set functions. *Mathematical Programming*, 14:265–294, 12 1978. doi: 10.1007/BF01588971.
- [14] Jure Leskovec, Andreas Krause, Carlos Guestrin, Christos Faloutsos, Jeanne VanBriesen, and Natalie Glance. Cost-effective outbreak detection in networks. In *Proceedings of the 13th ACM SIGKDD international conf. on Knowledge discovery and data mining*, pages 420–429, 2007.
- [15] Uriel Feige. A threshold of $\ln n$ for approximating set cover. *J. ACM*, 45(4):634–652, jul 1998. ISSN 0004-5411. doi: 10.1145/285055.285059. URL <https://doi.org/10.1145/285055.285059>.
- [16] Bernhard Schölkopf, Alexander J. Smola, and Francis Bach. *Learning with Kernels: Support Vector Machines, Regularization, Optimization, and Beyond*. The MIT Press, 2018. ISBN 0262536579.

- [17] Carl Edward Rasmussen and Christopher K. I. Williams. *Gaussian Processes for Machine Learning (Adaptive Computation and Machine Learning)*. The MIT Press, 2005. ISBN 026218253X.
- [18] Sayak Ray Chowdhury and Aditya Gopalan. On kernelized multi-armed bandits. In *International Conference on Machine Learning*, pages 844–853. PMLR, 2017.
- [19] Niranjan Srinivas, Andreas Krause, Sham M. Kakade, and Matthias W. Seeger. Information-theoretic regret bounds for gaussian process optimization in the bandit setting. *IEEE Transactions on Information Theory*, 58(5):3250–3265, 2012. doi: 10.1109/TIT.2011.2182033.
- [20] Johannes Kirschner, Tor Lattimore, and Andreas Krause. Information directed sampling for linear partial monitoring. In *Conference on Learning Theory*, pages 2328–2369. PMLR, 2020.
- [21] Sattar Vakili, Kia Khezeli, and Victor Picheny. On information gain and regret bounds in gaussian process bandits. In *International Conference on Artificial Intelligence and Statistics*, pages 82–90. PMLR, 2021.
- [22] Open weather. <https://openweathermap.org/>, 2022.
- [23] Mojmír Mutný and Andreas Krause. Sensing cox processes via posterior sampling and positive bases. *CoRR*, abs/2110.11181, 2021. URL <https://arxiv.org/abs/2110.11181>.
- [24] Neba Funwi-gabga and Jorge Mateu. Understanding the nesting spatial behaviour of gorillas in the kagwene sanctuary, cameroon. *Stochastic Environmental Research and Risk Assessment*, 26, 08 2011. doi: 10.1007/s00477-011-0541-1.
- [25] Akifumi Wachi and Yanan Sui. Safe reinforcement learning in constrained markov decision processes. In *ICML*, pages 9797–9806, 2020. URL <http://proceedings.mlr.press/v119/wachi20a.html>.
- [26] Yanan Sui, Alkis Gotovos, Joel Burdick, and Andreas Krause. Safe exploration for optimization with gaussian processes. In *International conference on machine learning*, pages 997–1005. PMLR, 2015.
- [27] Bobak Shahriari, Kevin Swersky, Ziyu Wang, Ryan P. Adams, and Nando de Freitas. Taking the human out of the loop: A review of bayesian optimization. *Proceedings of the IEEE*, 104 (1):148–175, 2016. doi: 10.1109/JPROC.2015.2494218.
- [28] Tor Lattimore and Csaba Szepesvári. *Partial Monitoring*, page 423–451. Cambridge University Press, 2020. doi: 10.1017/9781108571401.046.
- [29] Tor Lattimore and Csaba Szepesvári. An information-theoretic approach to minimax regret in partial monitoring. In *Conference on Learning Theory*, pages 2111–2139. PMLR, 2019.
- [30] Lai Wei, Andrew McDonald, and Vaibhav Srivastava. Regret analysis of distributed gaussian process estimation and coverage. *CoRR*, abs/2101.04306, 2021. URL <https://arxiv.org/abs/2101.04306>.
- [31] Andrea Carron, Marco Todescato, Ruggero Carli, Luca Schenato, and Gianluigi Pillonetto. Multi-agents adaptive estimation and coverage control using gaussian regression. In *2015 European Control Conference (ECC)*, pages 2490–2495, 2015. doi: 10.1109/ECC.2015.7330912.
- [32] Andreas Krause, Ajit Singh, and Carlos Guestrin. Near-optimal sensor placements in gaussian processes: Theory, efficient algorithms and empirical studies. *J. Mach. Learn. Res.*, 9:235–284, jun 2008. ISSN 1532-4435.
- [33] S. Lloyd. Least squares quantization in pcm. *IEEE Transactions on Information Theory*, 28(2): 129–137, 1982. doi: 10.1109/TIT.1982.1056489.
- [34] Jorge Cortes, Sonia Martinez, Timur Karatas, and Francesco Bullo. Coverage control for mobile sensing networks. *IEEE Transactions on robotics and Automation*, 20(2):243–255, 2004.

- [35] Francois Lekien and Naomi Ehrich Leonard. Nonuniform coverage and cartograms. *SIAM Journal on Control and Optimization*, 48(1):351–372, 2009. doi: 10.1137/070681120. URL <https://doi.org/10.1137/070681120>.
- [36] Yunfei Xu and Jongeun Choi. Adaptive sampling for learning gaussian processes using mobile sensor networks. *Sensors (Basel, Switzerland)*, 11:3051–66, 12 2011. doi: 10.3390/s110303051.
- [37] Wenhao Luo and Katia Sycara. Adaptive sampling and online learning in multi-robot sensor coverage with mixture of gaussian processes. In *2018 IEEE International Conference on Robotics and Automation (ICRA)*, pages 6359–6364. IEEE, 2018.
- [38] Alessia Benevento, María Santos, Giuseppe Notarstefano, Kamran Paynabar, Matthieu Bloch, and Magnus Egerstedt. Multi-robot coordination for estimation and coverage of unknown spatial fields. In *2020 IEEE International Conference on Robotics and Automation (ICRA)*, pages 7740–7746, 2020. doi: 10.1109/ICRA40945.2020.9197487.
- [39] Walker Gosrich, Siddharth Mayya, Rebecca Li, James Paulos, Mark Yim, Alejandro Ribeiro, and Vijay Kumar. Coverage control in multi-robot systems via graph neural networks. In *2022 International Conference on Robotics and Automation (ICRA)*, pages 8787–8793. IEEE, 2022.
- [40] Andrea Carron and Melanie N Zeilinger. Model predictive coverage control. *IFAC-PapersOnLine*, 53(2):6107–6112, 2020.
- [41] Soobum Kim, María Santos, Luis Guerrero-Bonilla, Anthony Yezzi, and Magnus Egerstedt. Coverage control of mobile robots with different maximum speeds for time-sensitive applications. *IEEE Robotics and Automation Letters*, 7(2):3001–3007, 2022. doi: 10.1109/LRA.2022.3146593.
- [42] Francesco Bullo, Ruggero Carli, and Paolo Frasca. Gossip coverage control for robotic networks: Dynamical systems on the space of partitions. *SIAM Journal on Control and Optimization*, 50(1):419–447, 2012.
- [43] Saba Faryadi and Javad Velni. A reinforcement learning-based approach for modeling and coverage of an unknown field using a team of autonomous ground vehicles. *International Journal of Intelligent Systems*, 36, 11 2020. doi: 10.1002/int.22331.
- [44] Gianpietro Battocletti, Riccardo Urban, Simone Godio, and Giorgio Guglieri. RL-based path planning for autonomous aerial vehicles in unknown environments. In *AIAA AVIATION 2021 FORUM*, page 3016, 2021.
- [45] Jeff Bilmes. Submodularity in machine learning and artificial intelligence. *arXiv preprint arXiv:2202.00132*, 2022.
- [46] Andreas Krause and Daniel Golovin. Submodular function maximization. *Tractability*, 3: 71–104, 2014.
- [47] Andreas Krause and Carlos Guestrin. Near-optimal nonmyopic value of information in graphical models. In *Proceedings of the Twenty-First Conference on Uncertainty in Artificial Intelligence*, UAI’05, page 324–331. AUAI Press, 2005. ISBN 0974903914.
- [48] Lin Chen, Andreas Krause, and Amin Karbasi. Interactive submodular bandit. In *NIPS*, 2017.
- [49] Matthew Streeter, Daniel Golovin, and Stephen F Smith. Combining multiple heuristics online. In *AAAI*, pages 1197–1203, 2007.
- [50] Daniel Golovin, Andreas Krause, and Matthew Streeter. Online submodular maximization under a matroid constraint with application to learning assignments. *arXiv preprint arXiv:1407.1082*, 2014.
- [51] Adish Singla, Sebastian Tschiatschek, and Andreas Krause. Noisy submodular maximization via adaptive sampling with applications to crowdsourced image collection summarization. In *Thirtieth AAAI Conference on Artificial Intelligence*, 2016.
- [52] Yisong Yue and Carlos Guestrin. Linear submodular bandits and their application to diversified retrieval. *Advances in Neural Information Processing Systems*, 24, 2011.

- [53] Lukas Hewing, Kim P Wabersich, Marcel Menner, and Melanie N Zeilinger. Learning-based model predictive control: Toward safe learning in control. *Annual Review of Control, Robotics, and Autonomous Systems*, 3:269–296, 2020.
- [54] Matteo Turchetta, Andrey Kolobov, Shital Shah, Andreas Krause, and Alekh Agarwal. Safe reinforcement learning via curriculum induction. *Advances in Neural Information Processing Systems*, 33:12151–12162, 2020.
- [55] Felix Berkenkamp, Matteo Turchetta, Angela Schoellig, and Andreas Krause. Safe model-based reinforcement learning with stability guarantees. *Advances in neural information processing systems*, 30, 2017.
- [56] Yinlam Chow, Ofir Nachum, Edgar Duenez-Guzman, and Mohammad Ghavamzadeh. A lyapunov-based approach to safe reinforcement learning. *Advances in neural information processing systems*, 31, 2018.
- [57] Jaime F Fisac, Anayo K Akametalu, Melanie N Zeilinger, Shahab Kaynama, Jeremy Gillula, and Claire J Tomlin. A general safety framework for learning-based control in uncertain robotic systems. *IEEE Transactions on Automatic Control*, 64(7):2737–2752, 2018.
- [58] Joshua Achiam, David Held, Aviv Tamar, and Pieter Abbeel. Constrained policy optimization. In *International conference on machine learning*, pages 22–31. PMLR, 2017.
- [59] Jeremy Coulson, John Lygeros, and Florian Dorfler. Distributionally robust chance constrained data-enabled predictive control. *IEEE Transactions on Automatic Control*, 2021.
- [60] Lukas Brunke, Melissa Greeff, Adam W Hall, Zhaocong Yuan, Siqi Zhou, Jacopo Panerati, and Angela P Schoellig. Safe learning in robotics: From learning-based control to safe reinforcement learning. *Annual Review of Control, Robotics, and Autonomous Systems*, 5, 2021.
- [61] Alex Ray, Joshua Achiam, and Dario Amodei. Benchmarking safe exploration in deep reinforcement learning. *arXiv preprint arXiv:1910.01708*, 7:1, 2019.
- [62] Jan Leike, Miljan Martic, Victoria Krakovna, Pedro A Ortega, Tom Everitt, Andrew Lefrancq, Laurent Orseau, and Shane Legg. Ai safety gridworlds. *arXiv preprint arXiv:1711.09883*, 2017.
- [63] Yanan Sui, Vincent Zhuang, Joel W. Burdick, and Yisong Yue. Stagewise safe bayesian optimization with gaussian processes, 2018. URL <https://arxiv.org/abs/1806.07555>.
- [64] Matteo Turchetta, Felix Berkenkamp, and Andreas Krause. Safe exploration in finite markov decision processes with gaussian processes. *Advances in Neural Information Processing Systems*, 29, 2016.
- [65] Zheqing Zhu, Erdem Bıyık, and Dorsa Sadigh. Multi-agent safe planning with gaussian processes. In *International Conference on Intelligent Robots and Systems (IROS)*, pages 6260–6267, 10 2020. doi: 10.1109/IROS45743.2020.9341169.
- [66] Mojmír Mutný and Andreas Krause. Experimental design for linear functionals in reproducing kernel hilbert spaces. *arXiv preprint arXiv:2205.13627*, 2022.
- [67] Maximilian Balandat, Brian Karrer, Daniel Jiang, Samuel Daulton, Ben Letham, Andrew G Wilson, and Eytan Bakshy. Botorch: a framework for efficient monte-carlo bayesian optimization. *Advances in neural information processing systems*, 33:21524–21538, 2020.
- [68] Jacob Gardner, Geoff Pleiss, Kilian Q Weinberger, David Bindel, and Andrew G Wilson. Gpytorch: Blackbox matrix-matrix gaussian process inference with gpu acceleration. *Advances in neural information processing systems*, 31, 2018.
- [69] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, et al. Pytorch: An imperative style, high-performance deep learning library. *Advances in neural information processing systems*, 32, 2019.

Checklist

1. For all authors...
 - (a) Do the main claims made in the abstract and introduction accurately reflect the paper’s contributions and scope? [\[Yes\]](#) See Section 5 for the main theorems and Section 6 for the experimental results
 - (b) Did you describe the limitations of your work? [\[Yes\]](#) See Section 8
 - (c) Did you discuss any potential negative societal impacts of your work? [\[N/A\]](#)
 - (d) Have you read the ethics review guidelines and ensured that your paper conforms to them? [\[Yes\]](#)
2. If you are including theoretical results...
 - (a) Did you state the full set of assumptions of all theoretical results? [\[Yes\]](#) See Section 3
 - (b) Did you include complete proofs of all theoretical results? [\[Yes\]](#) See Section 5 with corresponding links to the Appendix for Proofs
3. If you ran experiments...
 - (a) Did you include the code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL)? [\[Yes\]](#) Code is attached in the supplemental material along with the environment maps used. The code folder contains a ReadMe file containing instructions to reproduce the result.
 - (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)? [\[Yes\]](#) See Section 6 with corresponding links to Appendix G for complete experimental setup details
 - (c) Did you report error bars (e.g., with respect to the random seed after running experiments multiple times)? [\[Yes\]](#) Errors bars are reported by running experiments with multiple random seeds and random environments
 - (d) Did you include the total amount of compute and the type of resources used (e.g., type of GPUs, internal cluster, or cloud provider)? [\[Yes\]](#) Compute details are in Appendix G
4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets...
 - (a) If your work uses existing assets, did you cite the creators? [\[Yes\]](#) In Section 6 we cited the data and the model used in the Gorilla nest dataset
 - (b) Did you mention the license of the assets? [\[Yes\]](#) In Appendix G along with the experiment setup details we mentioned the license of the relevant code and data used
 - (c) Did you include any new assets either in the supplemental material or as a URL? [\[N/A\]](#)
 - (d) Did you discuss whether and how consent was obtained from people whose data you’re using/curating? [\[N/A\]](#)
 - (e) Did you discuss whether the data you are using/curating contains personally identifiable information or offensive content? [\[N/A\]](#)
5. If you used crowdsourcing or conducted research with human subjects...
 - (a) Did you include the full text of instructions given to participants and screenshots, if applicable? [\[N/A\]](#)
 - (b) Did you describe any potential participant risks, with links to Institutional Review Board (IRB) approvals, if applicable? [\[N/A\]](#)
 - (c) Did you include the estimated hourly wage paid to participants and the total amount spent on participant compensation? [\[N/A\]](#)

Part I

Appendix

Table of Contents

A	Definitions	17
A.1	Notations	17
A.2	GoOSE operators	19
A.3	Batching operation	20
B	Disk Coverage as a submodular function	20
C	Agent wise regret bound	22
C.1	Unconstrained case	22
C.2	Constrained case	25
D	Proof. for Theorem 1 (MACOPT)	28
D.1	Variants of MACOPT	32
E	Proof. for Theorem 2 (SAFEMAC)	33
E.1	Intermediate recommendation is near-optimal at SAFEMAC's convergence . . .	39
F	Multi-agent GoOSE version	41
G	Experiments	43
G.1	Scaling in terms of agents and domain size	44

A Definitions

A.1 Notations

Problem Formulation			
F	\triangleq Submodular function, $F : 2^V \rightarrow \mathbb{R}$	β_t^q	\triangleq Scaling, defined as per [18]
V	\triangleq Domain	L_q	\triangleq Lipschitz constant
v	\triangleq An element in the domain V	ϵ_q	\triangleq Statistical confidence up to which constraint function q is learnt
$F(\cdot; \cdot, \cdot)$	\triangleq Coverage objective defined in Eq. (1)	$d(v, z)$	\triangleq Distance metric
i	\triangleq Agent index	σ_q	\triangleq Standard deviation of constraint observations noise
ρ	\triangleq Density function, $\rho : V \rightarrow \mathbb{R}$	σ_t^q	\triangleq Posterior standard deviation of q GP
q	\triangleq Constraint function, $q : V \rightarrow \mathbb{R}$	B_q	\triangleq Norm bound of q , $\ q\ _{k^q} \leq B_q$
D^i	\triangleq Sensing region around agent i	η_q	\triangleq Noise in constraint observations
$D^{1:i}$	$\triangleq \cup_{j=1}^i D^j$, union of sensing regions of agents $1 : i$	l_t^ρ	\triangleq LCB of the density at time t
D^{i-}	$\triangleq D^i \setminus D^{1:i-1}$, region occupied by agent i , but not by $1 : i-1$ agents	u_t^ρ	\triangleq UCB of the density at time t
\tilde{D}^i	\triangleq Sensing region occupied by greedy optimal location of agent i	β_t^ρ	\triangleq Scaling, defined as per [18]
\tilde{D}^{i-}	$\triangleq \tilde{D}^i \setminus D^{1:i-1}$	w_t	$\triangleq \sum_{i=1}^N u_{t-1}^\rho(x_t^{g,i}) - l_{t-1}^\rho(x_t^{g,i})$, sum of highest uncertainty below disks
C_D	$\triangleq \max_i D^i / V \leq 1$, Maximum fraction of area covered by an agent	ϵ_ρ	\triangleq Accuracy threshold for learning the density, $w \leq \epsilon_\rho$
N	\triangleq Total number of agents	σ_t^ρ	\triangleq Posterior standard deviation of ρ GP
Batch Operation		σ_ρ	\triangleq Standard deviation of density observations noise
B	\triangleq A batch of agents, $\{1, 2 \dots B \}$	B_ρ	\triangleq Norm bound of ρ , $\ \rho\ _{k^\rho} \leq B_\rho$
$B'_t(i)$	$\triangleq \{j \in [N] S_t^{u,i} \cap S_t^{u,j} \neq \emptyset\}$, agents connected in union set with agent i	δ	$\triangleq \in (0, 1)$ for high probability
\mathcal{B}_t	$\triangleq \bigcup_{i \in [N]} B'_t(i)$. Collection of batches sharing the union set.	$H(y_A)$	\triangleq Shannon entropy
\mathcal{B}	\triangleq Collection of batches sharing the largest reachable set ($\bar{R}_{\epsilon_q}(X_0^B)$)	$I(y_A; \rho)$	$\triangleq H(y_A) - H(y_A \rho)$, Information gain
\mathcal{B}_t^p	\triangleq Collection of batches sharing the pessimistic set	γ	\triangleq Information capacity
X Notations		γ_{NT}^ρ	$\triangleq \sup_{A \subseteq V} I(Y_A; \rho)$, A is set of NT obs. $\gamma_{NT}^\rho := \gamma_{NT, \rho}$, ρ is clear in T .
x_t^i	\triangleq Planned location of agent i at time t	γ_{NT}^q	$\triangleq \sup_{A \subseteq V} I(Y_A; q)$, A is set of NT obs. $\gamma_{NT}^q := \gamma_{NT, q}$, q is clear in T .
$x_t^{g,i}$	\triangleq Goal of agent i at time t , defined by Line 6 and Line 8 in Algorithm 4	Tr	\triangleq Trace of a Matrix
\tilde{x}^i	\triangleq Greedy location of agent i , Eq. (19)	K^ρ	\triangleq Posterior kernel matrix with ρ obs.
X_t	$\triangleq \cup_{i \in [N]} \{x_t^i\}$, A set of agents at time t	$\lambda_{i,t}$	\triangleq Eigenvalue of the kernel matrix
X_t^B	$\triangleq \cup_{i \in B} \{x_t^i\}$, Agents in B at time t	η_ρ	\triangleq Noise in the density observations
X_\star^B	\triangleq Optimal location of batch B agents	Time	
X_\star	$\triangleq \cup_{B \in \mathcal{B}} X_\star^B$	t	\triangleq Any round of the algorithm
$X^{1:i}$	\triangleq A set of agents 1 to i	T	\triangleq Algorithm termination time
$x_{1:T}^{g,1:N}$	\triangleq A set of $1 : N$ agents' goal locations up to time T	t_q^\star	\triangleq Maximum number of q observations
Density (ρ) and Constraint (q) GP		t_ρ^\star	\triangleq Maximum number of ρ observations
l_t^q	\triangleq LCB of the constraint at time t	$t_\rho^{\star 1}$	\triangleq Maximum number of density observations for the first coverage phase
u_t^q	\triangleq UCB of the constraint at time t	$\delta t_\rho^{\star n}$	\triangleq Maximum number of density obs. from $(n-1)^{th}$ to n^{th} coverage phase
		δt_ρ^n	\triangleq Number of density obs. from $(n-1)^{th}$ to n^{th} coverage phase
		t_ρ^n	\triangleq Number of density observations till n^{th} coverage phase

GoOSE and Safe Expansion

$p_t(S)$	\triangleq pessimistic operator $\{v \in V, \exists z \in S : l_t^q(z) - L_q d(v, z) \geq 0\}$
$o_t^{\epsilon_q}(S)$	\triangleq optimistic operator $\{v \in V, \exists z \in S : u_t^q(z) - \epsilon_q - L_q d(v, z) \geq 0\}$
$\bar{P}_t(\cdot)$	\triangleq Pessimistic expansion operator
$\bar{O}_t(\cdot)$	\triangleq Optimistic expansion operator
$\bar{R}_{\epsilon_q}(\{x_0^i\})$	\triangleq Maximum safely reachable set up to ϵ_q , Eq. (12)
$\bar{R}_{\epsilon_q}(X_0^B)$	$\triangleq \cup_{i \in B} \bar{R}_{\epsilon_q}(\{x_0^i\})$
$S_t^{p,i}$	\triangleq Pessimistic set of agent i , $\tilde{P}_t(S_{t-1}^{p,i})$
$S_t^{p,B}$	$\triangleq \cup_{i \in B} S_t^{p,i}$
S_t^p	\triangleq Pessimistic set of all N agents
$S_t^{o,\epsilon_q,i}$	\triangleq Optimistic set of agent i , $\tilde{O}_t^{\epsilon_q}(S_{t-1}^{p,i})$
$S_t^{o,\epsilon_q,B}$	$\triangleq \cup_{i \in B} S_t^{o,\epsilon_q,i}$
S_t^{o,ϵ_q}	\triangleq Optimistic set of all N agents
$S_t^{u,i}$	\triangleq Union set, $S_t^{o,\epsilon_q,i} \cup S_t^{p,i}$
$S_t^{u,B}$	$\triangleq \cup_{i \in B} S_t^{u,i}$
S_t^u	\triangleq Union set of all N agents
$R_{\epsilon_q}^{\text{safe}}(S)$	\triangleq True safety constraint operator, Eq. (8)
$R_n^{\text{reach}}(S)$	\triangleq n step reachability in the graph, Eq. (9)
$\tilde{R}^{\text{reach}}(S)$	$\triangleq \lim_{n \rightarrow \infty} R_n^{\text{reach}}(S)$
$R_{\epsilon_q}^n(S)$	\triangleq n step safely reachable set in the graph, Eq. (12)
$\bar{R}_{\epsilon_q}(S)$	$\triangleq \lim_{n \rightarrow \infty} R_{\epsilon_q}^n(S)$
$W_t^{\epsilon_q}$	\triangleq Set of locations whose safety is not ϵ_q -accurate, Algorithm 3
$G_t^{\epsilon_q}(\alpha)$	\triangleq A set of potential immediate expanders, Algorithm 3
p	\triangleq Priority, Algorithm 3
$h(v)$	\triangleq Heuristic function, Algorithm 3
$A_t(\alpha)$	\triangleq Subset of locations with equal priority, Algorithm 3

Regret

$F(X)$	$\triangleq F(X; \rho, V)$, short notation when ρ and V are obvious
$\Delta(x^i X^{1:i-1}; \rho, V)$	\triangleq Marginal coverage gain by agent i , Eq. (17)
$\Delta(x^i X^{1:i-1})$	$\triangleq \Delta(x^i X^{1:i-1}; \rho, V)$, short notation when ρ and V are obvious
$Reg_{act}(T)$	\triangleq Actual regret in unconstrained case, Eq. (5)
OPT_l^i	\triangleq Per agent cumulative optimal gain, Eq. (21)
$Reg^i(T)$	\triangleq Per agent regret, Eq. (22)
OPT	$\triangleq \sum_{t=1}^T F(X_*)$
r_t^{act}	\triangleq Simple actual regret, constrained case, Eq. (29)
r_t^O	\triangleq Simple actual regret in union set, constrained case, Eq. (29)
r_t	\triangleq Simple per agent regret, constrained case, Eq. (29)
$Reg_{act}^O(T)$	\triangleq Cumulative actual regret, Eq. (30)
$Reg_l^O(T)$	\triangleq Sum of cumulative per agent regret, Eq. (30)

A.2 GoOSE operators

We denote with $\mathcal{G} = (V, \mathcal{E})$ the undirected graph describing the dependency among locations, V indicates the vertices of the graph, i.e., the state space of the problem and $\mathcal{E} \subseteq V \times V$ denotes the edges. In our setting, there are N identical agents having the same transition dynamics. Each agent can have a separate $\tilde{R}_{\epsilon_q}(\{x_0^i\})$.

The baseline as per true safety constraint operator:

$$R_{\epsilon_q}^{\text{safe}}(S) = S \cup \{v \in V \setminus S, |\exists z \in S : q(z) - \epsilon_q - L_q d(v, z) \geq 0\} \quad (8)$$

Now, we define reachability operator as all the locations that can be reached starting from set S .

$$R^{\text{reach}}(S) = S \cup \{v \in V \setminus S, |\exists z \in S : (z, v) \in \mathcal{E}\},$$

$$R_n^{\text{reach}}(S) = R_n^{\text{reach}}(R_{n-1}^{\text{reach}}(S)) \text{ with } R_1^{\text{reach}}(S) = R^{\text{reach}}(S) \quad (9)$$

$$\tilde{R}^{\text{reach}}(S) = \lim_{n \rightarrow \infty} R_n^{\text{reach}}(S), \quad (10)$$

For defining $\bar{R}_{\epsilon_q}(S)$,

$$R_{\epsilon_q}(S) = R_{\epsilon_q}^{\text{safe}}(S) \cap \tilde{R}^{\text{reach}}(S)$$

$$R_{\epsilon_q}^n(S) = R_{\epsilon_q}(R_{\epsilon_q}^{n-1}(S)) \text{ with } R_{\epsilon_q}^1(S) = R_{\epsilon_q}(S) \quad (11)$$

$$\bar{R}_{\epsilon_q}(S) = \lim_{n \rightarrow \infty} R_{\epsilon_q}^n(S) \quad (12)$$

Optimistic and pessimistic constrain satisfaction operators:

$$o_t^{\epsilon_q}(S) = \{v \in V, |\exists z \in S : u_t^q(z) - \epsilon_q - L_q d(v, z) \geq 0\}$$

$$p_t^{\epsilon_q}(S) = \{v \in V, |\exists z \in S : l_t^q(z) - \epsilon_q - L_q d(v, z) \geq 0\}$$

In this section, for simplicity, we have considered an undirected graph. This results in the same reachability and returnability operators since the edges are bidirectional. The extension to the directed graph is easy by using the reachability, the returnability and the ergodic operator. (Appendix A of Turchetta et al. [12] does it for the directed graph, so we did not repeat it here)

The optimistic and pessimistic expansion operators are given by,

$$O_t^{\epsilon_q}(S) = o_t^{\epsilon_q}(S) \cap \tilde{R}^{\text{reach}}(S)$$

$$O_t^{\epsilon_q, n}(S) = O_t^{\epsilon_q}(O_t^{\epsilon_q, n-1}(S)) \text{ with } O_t^{\epsilon_q, 1}(S) = O_t^{\epsilon_q}(S)$$

$$\tilde{O}_t^{\epsilon_q}(S) = \lim_{n \rightarrow \infty} O_t^{\epsilon_q, n}(S)$$

Pessimistic expansion operator

$$P_t^{\epsilon_q}(S) = p_t^{\epsilon_q}(S) \cap \tilde{R}^{\text{reach}}(S)$$

$$P_t^{\epsilon_q, n}(S) = P_t^{\epsilon_q}(P_t^{\epsilon_q, n-1}(S)) \text{ with } P_t^{\epsilon_q, 1}(S) = P_t^{\epsilon_q}(S)$$

$$\tilde{P}_t^{\epsilon_q}(S) = \lim_{n \rightarrow \infty} P_t^{\epsilon_q, n}(S)$$

This gives the optimistically and pessimistically, safe and reachable set respectively as:

$$S_t^{o, \epsilon_q} = \tilde{O}_t^{\epsilon_q}(S_{t-1}^p)$$

$$S_t^p = \tilde{P}_t^0(S_{t-1}^p)$$

Now in our setting with N agents, we denote with $S_t^{o, \epsilon_q, i}$ and $S_t^{p, i}$, the optimistic and the pessimistic set respectively of agent i . The union set for any agent i is defined as,

$$S_t^{u, i} := S_t^{o, \epsilon_q, i} \cup S_t^{p, i} \quad (13)$$

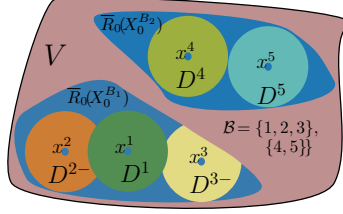


Figure 5: Disconnected safe regions: Agents are partitioned into two batches. Agent 1 covers D^1 (green), 2 covers D^{2-} (orange) and 3 covers D^{3-} (yellow).

A.3 Batching operation

For a set of agents, we partition them in batches, such that each batch B contains the agents that share at least a node in the union set. The total collection of batches, \mathcal{B} , is defined as,

$$\mathcal{B}_t = \bigcup_{i \in [N]} \mathcal{B}'_t(i) \text{ where } \mathcal{B}'_t(i) = \{j \in [N] \mid S_t^{u,i} \cap S_t^{u,j} \neq \emptyset\} \quad (14)$$

Analogous to \mathcal{B}_t , we define \mathcal{B}_t^p (or \mathcal{B}) as collection of batches where any $B \in \mathcal{B}_t^p$ (or \mathcal{B}) contains agents which are topologically connected in the pessimistic (or maximum safely reachable) set. Precisely,

$$\mathcal{B}_t^p = \bigcup_{i \in [N]} \mathcal{B}'_t(i) \text{ where } \mathcal{B}'_t(i) = \{j \in [N] \mid S_t^{p,i} \cap S_t^{p,j} \neq \emptyset\} \quad (15)$$

$$\mathcal{B} = \bigcup_{i \in [N]} \mathcal{B}'(i) \text{ where } \mathcal{B}'(i) = \{j \in [N] \mid \bar{R}_{\epsilon_q}(\{x_0^i\}) \cap \bar{R}_{\epsilon_q}(\{x_0^j\}) \neq \emptyset\} \quad (16)$$

The resulting batch collection are mutually exclusive that is $\forall B_1, B_2 \in \mathcal{B}_t, B_1 \neq B_2, B_1 \cap B_2 = \emptyset$ and also, $\sum_{B \in \mathcal{B}_t} |B| = N$.

For any batch B we can define their combined union set, pessimistic set and the maximum safely reachable set as ,

$$S_t^{u,B} := \cup_{i \in B} S_t^{u,i}, \quad S_t^{p,B} := \cup_{i \in B} S_t^{p,i}, \quad \bar{R}_{\epsilon_q}(X_0^B) = \cup_{i \in B} \bar{R}_{\epsilon_q}(\{x_0^i\}).$$

B Disk Coverage as a submodular function

Set functions Function $F : 2^V \rightarrow \mathbb{R}$ that assign each subset $A \subseteq V$ a value $F(A)$.

Discrete Derivative For a set function $F : 2^V \rightarrow \mathbb{R}$, $A \subseteq V$, and $e \in V$, let $\Delta_F(e|A) := F(A \cup \{e\}) - F(A)$ is discrete derivative of F at A with respect to e .

Submodular functions A function $F(\cdot)$ is a submodular if, $\forall A \subseteq B \subseteq V$ and $\forall e \in V \setminus B$

$$F(A \cup \{e\}) - F(A) \geq F(B \cup \{e\}) - F(B), \\ \Delta_F(e|A) \geq \Delta_F(e|B).$$

For the disk coverage function $F(A)$, defined in Eq. (1),

$$F(X; \rho, V) = \sum_{x^i \in X} \sum_{v \in D^{i-}} \rho(v)/|V|,$$

We can write marginal gain as,

$$\begin{aligned} F(A \cup \{e\}) - F(A) &= \sum_{x^i \in A \cup \{e\}} \sum_{v \in D^{i-}} \rho(v)/|V| - \sum_{x^i \in A} \sum_{v \in D^{i-}} \rho(v)/|V| \\ &= \sum_{x^i \in A} \sum_{v \in D^{i-}} \rho(v)/|V| + \sum_{x^i \in \{e\}} \sum_{v \in D^i \setminus D^{1:|A|}} \rho(v)/|V| - \sum_{x^i \in A} \sum_{v \in D^{i-}} \rho(v)/|V| \\ &= \sum_{x^i \in \{e\}} \sum_{v \in D^i \setminus D^{1:|A|}} \rho(v)/|V| \end{aligned}$$

$$\begin{aligned}
&\geq \sum_{x^i \in \{e\}} \sum_{v \in D^i \setminus D^{1:|B|}} \rho(v)/|V| \quad (\text{Since, } A \subseteq B, |D^i \setminus D^{1:|A|}| \geq |D^i \setminus D^{1:|B|}|) \\
&= \sum_{x^i \in B \cup \{e\}} \sum_{v \in D^{i-}} \rho(v)/|V| - \sum_{x^i \in B} \sum_{v \in D^{i-}} \rho(v)/|V| \\
&= F(B \cup \{e\}) - F(B) \\
\implies F(A \cup \{e\}) - F(A) &\geq F(B \cup \{e\}) - F(B)
\end{aligned}$$

This shows that the coverage function defined in Eq. (1) is a Submodular function.

Monotonicity is directly implied by the definition of $F(A)$, as an additive function of ρ . Since, $\rho(v) \geq 0, \forall v \in V \implies F(A) \leq F(B)$, if $A \subseteq B$.

C Agent wise regret bound

In this section, we upper bound the actual ("greedy") regret with the per agent regret in the unconstrained and the constrained case. The proof methodology to bound with per agent regret is motivated from [52]. We first define marginal gain and agent-wise regret. Then we give a proposition for the submodularity rate equation, which will be central to our lemmas. Finally, we bound the actual regret with the sum of per agent regret for unconstrained and then constrained case in

Marginal coverage gain:

$$\begin{aligned}
\Delta(x_t^i | X_t^{1:i-1}; \rho, V) &= F(X_t^{1:i-1} \cup \{x_t^i\}; \rho, V) - F(X_t^{1:i-1}; \rho, V) \\
&= \sum_{x_t^i \in X_t^{1:i}} \sum_{v \in D_t^{i-}} \rho(v)/|V| - \sum_{x_t^i \in X_t^{1:i-1}} \sum_{v \in D_t^{i-}} \rho(v)/|V| \\
&= \sum_{v \in D_t^{i-}} \rho(v)/|V|
\end{aligned} \tag{17}$$

Using, $X^{1:0} = \{\emptyset\}$, $F(X^{1:0}) = 0$, it follows that,

$$\sum_{i=1}^N \Delta(x_t^i | X_t^{1:i-1}; \rho, V) = F(X_t^{1:N}; \rho, V) \tag{18}$$

Tilde Notations:

$$\tilde{x}_t^i = \arg \max_{x_t^i} \Delta(x_t^i | X_t^{1:i-1}; \rho, V) \tag{19}$$

Proposition 1 (Eq. (3-7), [46], Submodular rate equation). *For a monotone Submodular function F the following holds,*

$$\max_{x^i} F(X^{1:i-1} \cup \{x^i\}) - F(X^{1:i-1}) \geq \frac{F(X_\star) - F(X^{1:i-1})}{N}, \tag{20}$$

where $X^{1:i}$ is the set of i agents being picked greedily and N is the number of agents in X_\star .

Proof. Let $X_\star = \{x_\star^1, \dots, x_\star^N\}$

$$\begin{aligned}
F(X_\star) &\leq F(X_\star \cup X^{1:i-1}) && \text{(With monotonicity of } F) \\
&= F(X^{1:i-1}) + \sum_{j=1}^N \Delta(x_\star^j | X^{1:i-1} \cup \{x_\star^1, \dots, x_\star^{j-1}\}) && \text{(Telescopic sum)} \\
&\leq F(X^{1:i-1}) + \sum_{x \in X_\star} \Delta(x | X^{1:i-1}) && \text{(Follows by Submodularity of } F) \\
&\leq F(X^{1:i-1}) + \sum_{x \in X_\star} (F(X^{1:i} \cup \{x\}) - F(X^{1:i-1} \cup \{x\})) \\
&\quad \text{(since, } x^i \text{ is added greedily to maximize } \Delta(x | X^{1:i-1})) \\
&\leq F(X^{1:i-1}) + N(F(X^{1:i}) - F(X^{1:i-1})) && \text{(} N \text{ agents in } X_\star) \\
&\implies \frac{F(X_\star) - F(X^{1:i-1})}{N} \leq F(X^{1:i}) - F(X^{1:i-1})
\end{aligned}$$

The proposition follows directly since x^i is added greedily to $X^{1:i-1}$. \square

C.1 Unconstrained case

Note that for unconstrained case domain V and utility ρ is obvious, so for convenience we use short hand notation, i.e, $F(\cdot; \rho, V) = F(\cdot)$ and $\Delta(\cdot; \rho, V) = \Delta(\cdot)$.

Locally optimal gain. Let us define OPT_l^i as the local optimal coverage gained by agent i , given all the locations of agents $1 : i - 1$, formally given by,

$$OPT_l^i = \sum_{t=1}^T \left(\max_{x_t^i} F(X_t^{1:i-1} \cup \{x_t^i\}) - F(X_t^{1:i-1}) \right) = \sum_{t=1}^T \Delta(\tilde{x}_t^i | X_t^{1:i-1}) \quad (21)$$

We denote with OPT , the optimal coverage, precisely $OPT = \sum_{t=1}^T F(X_\star)$.

Per agent regret. Let us define local regret, as the difference in coverage gain in picking state \tilde{x}_t^i vs the picked location x_t^i (this disparity is due to not knowing the actual density)

$$Reg^i(T) = \sum_{t=1}^T \Delta(\tilde{x}_t^i | X_t^{1:i-1}) - \sum_{t=1}^T \Delta(x_t^i | X_t^{1:i-1}) = OPT_l^i - \sum_{t=1}^T \Delta(x_t^i | X_t^{1:i-1}) \quad (22)$$

Actual regret. The actual regret is given by,

$$Reg_{act}(T) = \left(1 - \frac{1}{e}\right) \sum_{t=1}^T F(X_\star) - \sum_{t=1}^T F(X_t) = \left(1 - \frac{1}{e}\right) OPT - \sum_{t=1}^T F(X_t) \quad (23)$$

To prove. In this section we aim to show that actual regret bounded by sum of per agent regret, precisely,

$$\begin{aligned} Reg_{act}(T) &\leq \sum_{i=1}^N Reg^i(T) \\ \sum_{i=1}^N Reg^i(T) &\geq \left(1 - \frac{1}{e}\right) OPT - \sum_{t=1}^T F(X_t^{1:N}) \quad (\text{Using def. of } Reg_{act}(T) \text{ from Eq. (23)}) \end{aligned}$$

Lemma 1. For all N agents' local per agent regret $Reg^i(T)$, we have,

$$\sum_{t=1}^T \Delta(x_t^i | X_t^{1:i-1}) \geq \frac{1}{N} \left(OPT - \sum_{t=1}^T F(X_t^{1:i-1}) \right) - Reg^i(T) \quad (24)$$

Proof.

$$\begin{aligned} \Delta(\tilde{x}_t^i | X_t^{1:i-1}) &= \max_{x_t^i} F(X_t^{1:i-1} \cup \{x_t^i\}) - F(X_t^{1:i-1}) && (\text{Using definition}) \\ &\geq \frac{F(X_\star) - F(X_t^{1:i-1})}{N} && (\text{Using Eq. (20) from Proposition 1}) \\ OPT_l^i &\geq \frac{1}{N} \left(\sum_{t=1}^T F(X_\star) - \sum_{t=1}^T F(X_t^{1:i-1}) \right) && (\text{Sum over time}) \\ &= \frac{1}{N} \left(OPT - \sum_{t=1}^T F(X_t^{1:i-1}) \right) && (\text{Using definition of } OPT) \\ \sum_{t=1}^T \Delta(x_t^i | X_t^{1:i-1}) &\geq \frac{1}{N} \left(OPT - \sum_{t=1}^T F(X_t^{1:i-1}) \right) - Reg^i(T) && (\text{Using def. of } Reg^i(T) \text{ Eq. (22)}) \end{aligned}$$

□

Lemma 2. For any time t , X_t being the recommended location by MACOPT, we have

$$\sum_{t=1}^T F(X_t^{1:N}) \geq \left(1 - \frac{1}{e}\right) OPT - \sum_{i=1}^N Reg^i(T) \quad (25)$$

And using definition of $Reg_{act}(T)$ from Eq. (23), this further implies that,

$$Reg_{act}(T) \leq \sum_{i=1}^N Reg^i(T) \quad (26)$$

Proof. The proof is similar to the Lemma 2 from [52]. We begin to prove by induction,

$$OPT - \sum_{t=1}^T F(X_t^{1:i}) \leq \left(1 - \frac{1}{N}\right)^i OPT + \sum_{m=1}^i Reg_l^m(T) \quad (27)$$

Our main goal, i.e., Eq. (25) can be proved by substituting $i = N$ and using the inequality $(1 - 1/N)^N < 1/e$ in Eq. (27).

For $i = 0$, corresponds to no agent case. So it's trivial.

Let's consider gap to optimal value, when i elements are already selected,

$$\begin{aligned} \delta^i &= OPT - \sum_{t=1}^T F(X_t^{1:i}) && \text{(LHS of Eq. (27))} \\ &= OPT - \sum_{t=1}^T \sum_{m=1}^i \Delta(x_t^m | X_t^{1:m-1}) && \text{(Sum marginal gain; Using Eq. (18))} \\ \delta^{i-1} &= OPT - \sum_{t=1}^T \sum_{m=1}^{i-1} \Delta(x_t^m | X_t^{1:m-1}) \\ \implies \delta^i &= \delta^{i-1} - \sum_{t=1}^T \Delta(x_t^i | X_t^{1:i-1}) && \text{(Subtract } \delta^{i-1} \text{ from } \delta^i) \\ \implies \sum_{t=1}^T \Delta(x_t^i | X_t^{1:i-1}) &= \delta^{i-1} - \delta^i && (28) \end{aligned}$$

This says that the gap to optimal reduces by $\sum_{t=1}^T \Delta(x_t^i | X_t^{1:i-1})$ after adding element $x_t^i \forall t$.

$$\begin{aligned} \sum_{t=1}^T \Delta(x_t^i | X_t^{1:i-1}) &\geq \frac{1}{N}(\delta^{i-1}) - Reg^i(T) && \text{(From Eq. (24) and } \delta^i \text{ definition)} \\ \implies \delta^{i-1} - \delta^i &\geq \frac{1}{N}(\delta^{i-1}) - Reg^i(T) && \text{(From Eq. (28))} \\ \implies \delta^i &\leq \left(1 - \frac{1}{N}\right)\delta^{i-1} + Reg^i(T) \\ &\leq \left(1 - \frac{1}{N}\right)^2 \delta^{i-2} + \sum_{m=1}^2 Reg^i(T) && \text{(Subs } \delta^{i-1}, \text{ Doing the telescopic bound)} \\ &\vdots \\ &\leq \left(1 - \frac{1}{N}\right)^i \delta^0 + \sum_{m=1}^i Reg^i(T) \\ &= \left(1 - \frac{1}{N}\right)^i OPT + \sum_{m=1}^i Reg^i(T) \\ OPT - \sum_{t=1}^T F(X_t^{1:i}) &\leq \left(1 - \frac{1}{N}\right)^i OPT + \sum_{m=1}^i Reg_l^m(T) && \text{(Using } \delta^i \text{ definition)} \end{aligned}$$

Hence proved. \square

C.2 Constrained case

Simple regret. We define for a particular t , simple regret r_t^{act} and per agent local regret r_t respectively as:

$$\begin{aligned} r_t^{act} &= (1 - \frac{1}{e}) \sum_{B \in \mathcal{B}} F(X_\star^B; \rho, \bar{R}_{\epsilon_q}(X_0^B)) - \sum_{B \in \mathcal{B}_t} \sum_{i \in B} \Delta(x_t^i | X_t^{1:i-1}; \rho, S_t^{u,B}), \\ r_t^O &= (1 - \frac{1}{e}) \sum_{B \in \mathcal{B}_t} F(X_\star^B; \rho, S_t^{u,B}) - \sum_{B \in \mathcal{B}_t} \sum_{i \in B} \Delta(x_t^i | X_t^{1:i-1}; \rho, S_t^{u,B}) \\ r_t &= \sum_{B \in \mathcal{B}_t} \sum_{i \in B} \Delta(\tilde{x}^i | X_t^{1:i-1}; \rho, S_t^{u,B}) - \Delta(x_t^i | X_t^{1:i-1}; \rho, S_t^{u,B}) \end{aligned} \quad (29)$$

Cumulative regret. The actual cumulative regret $Reg_{act}^O(T)$ and the per agent cumulative regret $Reg_l^O(T)$ are respectively given by,

$$Reg_{act}^O(T) = \sum_{t=1}^T r_t^{act} \text{ and } Reg_l^O(T) = \sum_{t=1}^T r_t \quad (30)$$

On bounding per batch regret.

Optimal coverage in a batch B

$$\begin{aligned} OPT_t &= F(X_\star^B; \rho, S_t^{u,B}) \\ OPT_t^i &= \max_{x^i} F(X_t^{1:i-1} \cup \{x^i\}; \rho, S_t^{u,B}) - F(X_t^{1:i-1}; \rho, S_t^{u,B}) \\ &= \max_{x^i} \Delta(x^i | X_t^{1:i-1}; \rho, S_t^{u,B}) = \Delta(\tilde{x}^i | X_t^{1:i-1}; \rho, S_t^{u,B}) \\ r_B^i(t) &= \Delta(\tilde{x}^i | X_t^{1:i-1}; \rho, S_t^{u,B}) - \Delta(x_t^i | X_t^{1:i-1}; \rho, S_t^{u,B}) \end{aligned} \quad (31)$$

To prove:

$$F(X_t^B; \rho, S_t^{u,B}) \geq (1 - \frac{1}{e}) OPT_t - \sum_{i \in B} r_B^i(t) \quad (32)$$

Proposition 2. Let N_B be the number of agents in batch B and for all such agents per agent regret is $r_B^i(t)$. Then the following holds,

$$\Delta(x_t^i | X_t^{1:i-1}; \rho, S_t^{u,B}) \geq \frac{1}{N_B} \left(OPT_t - F(X_t^{1:i-1}; \rho, S_t^{u,B}) \right) - r_B^i(t) \quad (33)$$

Proof.

$$\Delta(\tilde{x}_t^i | X_t^{1:i-1}; \rho, S_t^{u,B}) = \max_{x_t^i} F(X_t^{1:i-1} \cup \{x_t^i\}; \rho, S_t^{u,B}) - F(X_t^{1:i-1}; \rho, S_t^{u,B}) \quad (\text{Using definition})$$

$$\geq \frac{F(X_\star; \rho, S_t^{u,B}) - F(X_t^{1:i-1}; \rho, S_t^{u,B})}{N_B} \quad (\text{Using Eq. (20) from Proposition 1})$$

$$OPT_t^i \geq \frac{1}{N_B} \left(OPT_t - F(X_t^{1:i-1}; \rho, S_t^{u,B}) \right) \quad (\text{Using definition of } OPT_t \text{ and } OPT_t^i)$$

$$\Delta(x_t^i | X_t^{1:i-1}; \rho, S_t^{u,B}) \geq \frac{1}{N_B} \left(OPT_t - F(X_t^{1:i-1}; \rho, S_t^{u,B}) \right) - r_B^i(t) \quad (\text{Using def. of } r_B^i(t) \text{ Eq. (31)})$$

□

Lemma 3. For any time t , X_t^B being the recommended location by SAFEMAC in the union set $S_t^{u,B}$, we have

$$F(X_t^B; \rho, S_t^{u,B}) \geq \left(1 - \frac{1}{e}\right) OPT_t - \sum_{i \in B} r_B^i(t), \quad (34)$$

Proof. The proof is similar to the Lemma 2 from [52]. We begin to prove by induction,

$$OPT_t - F(X_t^{1:i}; \rho, S_t^{u,B}) \leq \left(1 - \frac{1}{N_B}\right)^i OPT_t + \sum_{m=1}^i r_B^m(t) \quad (35)$$

For $i = 0$, corresponds to no agent case. So it's trivial.

Let's consider gap to optimal value, when i elements are already selected,

$$\begin{aligned} \delta^i &= OPT_t - F(X_t^{1:i}; \rho, S_t^{u,B}) && \text{(LHS of Eq. (35))} \\ &= OPT_t - \sum_{m=1}^i \Delta(x_t^m | X_t^{1:m-1}; \rho, S_t^{u,B}) && \text{(sum of marginal gain)} \\ \delta^{i-1} &= OPT_t - \sum_{m=1}^{i-1} \Delta(x_t^m | X_t^{1:m-1}; \rho, S_t^{u,B}) \\ \implies \delta^i &= \delta^{i-1} - \Delta(x_t^i | X_t^{1:i-1}; \rho, S_t^{u,B}) && \text{(Subtract } \delta^{i-1} \text{ from } \delta^i) \\ \implies \Delta(x_t^i | X_t^{1:i-1}; \rho, S_t^{u,B}) &= \delta^{i-1} - \delta^i && (36) \end{aligned}$$

This says that the gap to optimal reduces by $\Delta(x_t^i | X_t^{1:i-1}; \rho, S_t^{u,B})$ after adding element x_t^i .

$$\begin{aligned} \Delta(x_t^i | X_t^{1:i-1}; \rho, S_t^{u,B}) &\geq \frac{1}{N_B} (\delta^{i-1} - r_B^i(t)) && \text{(From Eq. (33) and } \delta^i \text{ definition)} \\ \implies \delta^{i-1} - \delta^i &\geq \frac{1}{N_B} (\delta^{i-1} - r_B^i(t)) && \text{(From Eq. (28))} \\ \implies \delta^i &\leq \left(1 - \frac{1}{N_B}\right) \delta^{i-1} + r_B^i(t) \\ &\leq \left(1 - \frac{1}{N_B}\right)^2 \delta^{i-2} + \sum_{m=1}^2 r_B^m(t) && \text{(Subs } \delta^{i-1}, \text{ Doing the telescopic bound)} \\ &\vdots \\ &\leq \left(1 - \frac{1}{N_B}\right)^i \delta^0 + \sum_{m=1}^i r_B^m(t) \\ &= \left(1 - \frac{1}{N_B}\right)^i OPT_t + \sum_{m=1}^i r_B^m(t) \\ OPT_t - F(X_t^{1:i}; \rho, S_t^{u,B}) &\leq \left(1 - \frac{1}{N_B}\right)^i OPT_t + \sum_{m=1}^i r_B^m(t) && \text{(Using } \delta^i \text{ definition)} \end{aligned}$$

Our main goal, i.e., Eq. (34) can be proved by substituting $i = N$ and using the inequality $(1 - 1/N)^N < 1/e$ in Eq. (35). Hence proved. \square

On combining all the batches.

Lemma 4. For any time t , X_t being the location recommended by SAFEMAC, we have

$$r_t^{act} \leq r_t^O \leq r_t \quad (37)$$

This further implies that,

$$Reg_{act}^O(T) \leq Reg_l^O(T) \quad (38)$$

Proof. For a batch B of agents, using Eq. (35) from Lemma 3 and substituting $r_B^i(t)$ from Eq. (31) we know that,

$$\begin{aligned} (1 - \frac{1}{e})F(X_\star^B; \rho, S_t^{u,B}) - \sum_{i \in B} \Delta(x_t^i | X_t^{1:i-1}; \rho, S_t^{u,B}) \\ \leq \sum_{i \in B} \Delta(\tilde{x}^i | X_t^{1:i-1}; \rho, S_t^{u,B}) - \Delta(x_t^i | X_t^{1:i-1}; \rho, S_t^{u,B}) \end{aligned}$$

By summing over all the $B \in \mathcal{B}_t$, we get

$$\begin{aligned} r_t^O = (1 - \frac{1}{e}) \sum_{B \in \mathcal{B}_t} F(X_\star^B; \rho, S_t^{u,B}) - \sum_{B \in \mathcal{B}_t} \sum_{i \in B} \Delta(x_t^i | X_t^{1:i-1}; \rho, S_t^{u,B}) \\ \leq \sum_{B \in \mathcal{B}_t} \sum_{i \in B} \Delta(\tilde{x}^i | X_t^{1:i-1}; \rho, S_t^{u,B}) - \Delta(x_t^i | X_t^{1:i-1}; \rho, S_t^{u,B}) \quad (39) \end{aligned}$$

Note that in Eq. (29), both the X_\star^B represents optimal agent's location in their respective coverage set i.e, $\bar{R}_{\epsilon_q}(x_0^i)$ and $S_t^{u,B}$, hence both the X_\star^B are different. Since, $\bigcup_{i \in B} \bar{R}_{\epsilon_q}(\{x_0^i\}) \subseteq S_t^{o, \epsilon_q, B} \subseteq S_t^{u,B} \implies \sum_{B \in \mathcal{B}} F(X_\star^B; \rho, \bar{R}_{\epsilon_q}(X_0^B)) \leq \sum_{B \in \mathcal{B}_t} F(X_\star^B; \rho, S_t^{u,B})$,

Moreover on using Eq. (29), Eq. (39) and we can conclude,

$$r_t^{act} \leq r_t^O \leq r_t.$$

This further implies Eq. (38) using definition in Eq. (30). Hence Proved \square

D Proof. for Theorem 1 (MACOPT)

Theorem 1. Let $\delta \in (0, 1)$, $\beta_t^{\rho^{1/2}} = B_\rho + 4\sigma_\rho \sqrt{\gamma_{N_t}^\rho + \ln(1/\delta)}$ and $C_D = \max_{x^i \in V} |D^i|/|V| \leq 1$. With probability at least $1 - \delta$, MACOPT's regret defined in Eq. (5) is bounded by $\mathcal{O}(\sqrt{T\beta_T^\rho \gamma_{NT}^\rho})$,

$$\Pr \left\{ \text{Reg}_{act}(T) \leq \sqrt{\frac{8C_D N T \beta_T^\rho \gamma_{NT}^\rho}{\log(1 + N\sigma_\rho^{-2})}} \right\} \geq 1 - \delta. \quad (6)$$

Proof. The proof for Theorem 1 goes in the following steps:

1. We first exploit the conditional linearity of the submodular objective to bound the cumulative regret defined in Eq. (5) with a sum of per agent regret ($\sum_{i=1}^N \text{Reg}^i(T)$). Precisely, we show $\text{Reg}_{act}(T) \leq \sum_{i=1}^N \text{Reg}^i(T)$ in Lemma 2.
2. We next bound the per agent regret with the information capacity γ_{NT}^ρ , a quantity that measures the largest reduction in uncertainty about the density that can be obtained from NT noisy evaluations of it.
 - For this, we quantify the information MACOPT acquires through the noisy density observations in Lemma 5, through the information gain $I(y_A; \rho) = H(y_A) - H(y_A|\rho)$, where H denotes the Shannon entropy and A is the set of locations evaluated by MACOPT.
 - Next we bound the per agent regret $\text{Reg}^i(T)$ with the information gain in Lemma 6-7 which is in turn bounded by the maximum information capacity.

Finally, Theorem 1 is a direct consequence of Lemma 2 and Lemma 7. \square

In the end of the section, we proof Corollary 1 which guarantees near optimal result in finite time.

Lemma 5. The information gain for the points observed by MACOPT can be expressed as:

$$I(Y_{x_{1:T}^{g,1:N}}; \rho) = \frac{1}{2} \sum_{t=1}^T \log(\det(I + \sigma_\rho^{-2} K_{x_t^{g,1:N}}^\rho)) = \frac{1}{2} \sum_{t=1}^T \sum_{i=1}^N \log(1 + \sigma_\rho^{-2} \lambda_{i,t}),$$

where $x_{1:T}^{g,1:N}$ is the set of goal locations set by MACOPT for all $1 : N$ agents up to time T . $K_{x_t^{g,1:N}}^\rho$ is the positive definite kernel matrix formed by the observed locations and $\lambda_{i,t}$ represents eigenvalue of the matrix.

Proof. We can precisely quantify this notion through the information gain

$$I(Y_{x_{1:T}^{g,1:N}}; \rho) = H(Y_{x_{1:T}^{g,1:N}}) - H(Y_{x_{1:T}^{g,1:N}}|\rho) \quad (40)$$

where H denotes the Shannon entropy. It can be defined as,

$$H(Y_{x_{1:T}^{g,1:N}}) = H(Y_T^{1:N} | Y_{x_{1:T-1}^{g,1:N}}) + H(Y_{x_{1:T-1}^{g,1:N}}) \quad (\text{Defined } Y_T^{1:N} := \{y_T^1, y_T^2, \dots, y_T^N\})$$

$$= \frac{1}{2} \log(\det(2\pi e(\sigma^2 I + K_{x_T^{g,1:N}}^\rho))) + H(Y_{T-1}^{1:N} | Y_{x_{1:T-2}^{g,1:N}}) + \dots \quad (41)$$

$$= \frac{1}{2} N \log(2\pi e \sigma^2) + \frac{1}{2} \log(\det(I + \sigma_\rho^{-2} K_{x_T^{g,1:N}}^\rho)) + H(Y_{T-1}^{1:N} | Y_{x_{1:T-2}^{g,1:N}}) + \dots \quad (42)$$

$$= \frac{1}{2} \sum_{t=1}^T N \log(2\pi e \sigma^2) + \frac{1}{2} \sum_{t=1}^T \log(\det(I + \sigma_\rho^{-2} K_{x_t^{g,1:N}}^\rho)) \quad (43)$$

For Eq. (41), we used that, $Y_T^{1:N} \sim \mathcal{N}(\mu_{T-1}^\rho(x_T^{g,1:N}), \sigma^2 I + K_{x_T^{g,1:N}}^\rho)$ is jointly a multivariate Gaussian. Eq. (42) follows by simplifying det, precisely, $\frac{1}{2} \log(\det(2\pi e(\sigma^2 I + K_{x_T^{g,1:N}}^\rho))) = \frac{1}{2} \log((2\pi e \sigma^2)^N \det(I + \sigma_\rho^{-2} K_{x_T^{g,1:N}}^\rho))$ and finally Eq. (43) by recursively repeating above 2 steps till

$t = 1$. $H(Y_{x_{1:T}^{g,1:N}}|\rho) = \frac{1}{2} \sum_{t=1}^T N \log(2\pi e \sigma^2)$ is the entropy because of the noise. On substituting this, with Eq. (43) in Eq. (40) we obtain,

$$\begin{aligned} I(Y_{x_{1:T}^{g,1:N}}; \rho) &= \frac{1}{2} \sum_{t=1}^T \log(\det(I + \sigma_\rho^{-2} K_{x_t^{g,1:N}}^\rho)) \\ &= \frac{1}{2} \sum_{t=1}^T \log\left(\prod_{i=1}^N (1 + \sigma_\rho^{-2} \lambda_{i,t})\right) \quad (\text{Using Eq. 45}) \\ &= \frac{1}{2} \sum_{t=1}^T \sum_{i=1}^N \log(1 + \sigma_\rho^{-2} \lambda_{i,t}) \end{aligned} \quad (44)$$

Hence Proved. \square

Log mat inequality:

$$\begin{aligned} \log(\det(I + \sigma_\rho^{-2} K^\rho)) &= \log(\det(RR^\top + \sigma_\rho^{-2} R\Lambda R^\top)) \quad (K^\rho = R\Lambda R^\top, RR^\top = I) \\ &= \log(\det(R(I + \sigma_\rho^{-2} \Lambda)R^\top)) \\ &= \log(\det(RR^\top)) + \log(\det(I + \sigma_\rho^{-2} \Lambda)) \quad (k \text{ is dimension of } K^\rho) \\ &= \log\left(\prod_{i=1}^k (1 + \sigma_\rho^{-2} \lambda_i)\right) \end{aligned} \quad (45)$$

Lemma 6. Till any time T , $\beta_t^{1/2} = B_\rho + 4\sigma_\rho \sqrt{\gamma_{N_t}^\rho + 1 + \ln(1/\delta)}$, if $|\rho(v) - \mu_{t-1}^\rho(v)| \leq \beta_t^{1/2} \sigma_{t-1}^\rho(v)$ for all $v \in V$, then the agent wise cumulative regret $\text{Reg}^i(T)$, is bounded by $\sum_{t=1}^T 2\sqrt{\beta_t^\rho} \sum_{v \in D_t^{i-}} \sigma_{t-1}^\rho(v)/|V|$ for agent i .

Proof. For notation convenience: $D_t^{i-} := D_t^i \setminus D_t^{1:i-1}$ and $\tilde{D}_t^{i-} := \tilde{D}_t^i \setminus D_t^{1:i-1}$

In MACOPT x_t^i is defined such that,

$$x_t^i = \arg \max_v \sum_{v \in D_t^{i-}} \mu_{t-1}^\rho(v) + \sqrt{\beta_t^\rho} \sigma_{t-1}^\rho(v) \quad (46)$$

Due to our picking strategy,

$$\sum_{v \in \tilde{D}_t^{i-}} \rho(v) \leq \sum_{v \in \tilde{D}_t^{i-}} (\mu_{t-1}^\rho(v) + \sqrt{\beta_t^\rho} \sigma_{t-1}^\rho(v)) \leq \sum_{v \in D_t^{i-}} (\mu_{t-1}^\rho(v) + \sqrt{\beta_t^\rho} \sigma_{t-1}^\rho(v)) \quad (47)$$

This first inequality follows due to upper bound and the second one follows based on how x_t^i is picked (Eq. (46)).

$$\begin{aligned} \text{Reg}^i(T) &= \sum_{t=1}^T \Delta(\tilde{x}_t^i | X_t^{1:i-1}) - \sum_{t=1}^T \Delta(x_t^i | X_t^{1:i-1}) \quad (\text{with definition Eq. (22)}) \\ &= \sum_{t=1}^T \left(\sum_{v \in \tilde{D}_t^{i-}} \rho(v) - \sum_{v \in D_t^{i-}} \rho(v) \right) / |V| \quad (\text{Using def. } \Delta(\cdot | X_t^{1:i-1}) \text{ Eq. (17)}) \\ &\leq \sum_{t=1}^T \left(\sum_{v \in D_t^{i-}} \mu_{t-1}^\rho(v) + \sqrt{\beta_t^\rho} \sigma_{t-1}^\rho(v) - \sum_{v \in D_t^{i-}} \rho(v) \right) / |V| \quad (\text{From Eq. (47)}) \\ &\leq \sum_{t=1}^T \left(\sum_{v \in D_t^{i-}} \mu_{t-1}^\rho(v) + \sqrt{\beta_t^\rho} \sigma_{t-1}^\rho(v) - \sum_{v \in D_t^{i-}} \mu_{t-1}^\rho(v) - \sqrt{\beta_t^\rho} \sigma_{t-1}^\rho(v) \right) / |V| \\ &\quad (\text{Since, } \rho(v) \geq \mu_{t-1}^\rho(v) - \sqrt{\beta_t^\rho} \sigma_{t-1}^\rho(v) \forall v) \end{aligned}$$

$$= \sum_{t=1}^T 2\sqrt{\beta_t^\rho} \sum_{v \in D_t^{i-}} \sigma_{t-1}^\rho(v)/|V| \quad (48)$$

□

Lemma 7. Let $\delta \in (0, 1)$ and let $\beta_t^{\rho^{1/2}} = B_\rho + 4\sigma_\rho \sqrt{\gamma_{Nt}^\rho + 1 + \ln(1/\delta)}$. Then for N agents, $\forall T \geq 1$ the following holds with probability $1 - \delta$,

$$\left(\sum_{i=1}^N \text{Reg}^i(T) \right)^2 \leq \frac{8C_D N T \beta_T^\rho I(Y_{x_{1:T}^{g,1:N}}; \rho)}{\log(1 + N\sigma_\rho^{-2})} \leq \frac{8C_D N T \beta_T^\rho \gamma_{NT}^\rho}{\log(1 + N\sigma_\rho^{-2})}$$

Proof. By sum over all the N agents from Lemma 6, we get

$$\sum_{i=1}^N \text{Reg}^i(T) \leq \sum_{i=1}^N \sum_{t=1}^T 2\sqrt{\beta_t^\rho} \sum_{v \in D_t^{i-}} \sigma_{t-1}^\rho(v)/|V| \quad (49)$$

Let's consider, part of Eq. (49), that is

$$\left(2\sqrt{\beta_t^\rho} \sum_{i=1}^N \sum_{v \in D_t^{i-}} \sigma_{t-1}^\rho(v)/|V| \right)^2 \leq 4\beta_t^\rho \sum_{i=1}^N \sum_{v \in D_t^{i-}} \left(\sigma_{t-1}^\rho(v) \right)^2 / |V| \quad (50)$$

$$\leq 4\beta_t^\rho \sum_{i=1}^N |D_t^{i-}| \left(\arg \max_{D^{i-}} \sigma_{t-1}^\rho(v) \right)^2 / |V| \quad (51)$$

$$\leq 4\beta_t^\rho C_D \sum_{i=1}^N \left(\sigma_{t-1}^\rho(x_t^{g,i}) \right)^2 \quad (52)$$

$$= 4\beta_t^\rho C_D \sum_{i=1}^N \lambda_{i,t} = 4\beta_t^\rho C_D \sum_{i=1}^N \sigma_\rho^2 \sigma_\rho^{-2} \lambda_{i,t} \quad (53)$$

$$\leq 4\beta_t^\rho C_D \sum_{i=1}^N \sigma_\rho^2 C_1 \log(1 + \sigma_\rho^{-2} \lambda_{i,t}) \quad (54)$$

$$\leq \frac{8C_D N \beta_t^\rho}{\log(1 + N\sigma_\rho^{-2})} \sum_{i=1}^N \frac{1}{2} \log(1 + \sigma_\rho^{-2} \lambda_{i,t}) \quad (55)$$

Eq. (50) follows from Cauchy-Schwarz inequality and using that $\sum_{i=1}^N |D^{i-}| \leq |V|$. Eq. (51) follows since $\sum_{v \in D_t^{i-}} (\sigma_{t-1}^\rho(v))^2 \leq |D^{i-}| \max_{v \in D_t^{i-}} (\sigma_{t-1}^\rho(v))^2$. We define $C_D = \frac{\max_i |D^i|}{|V|}$ denoting maximum coverage fraction possible by a disk. Eq. (53) follows from $\sum_{i=1}^N (\sigma_{t-1}^\rho(x_t^{g,i}))^2 = \text{Tr}(K^\rho) = \sum_{i=1}^N \lambda_{i,t}$. Since, $s \leq C_1 \log(1 + s)$ for $s \in [0, N\sigma_\rho^{-2}]$, where $C_1 = N\sigma_\rho^{-2} / \log(1 + N\sigma_\rho^{-2}) \geq 1$. Eq. (54) follows for $s = \sigma_\rho^{-2} \lambda_{i,t} \leq \sigma_\rho^{-2} \lambda_{\max} \leq \sigma_\rho^{-2} \sum_i \lambda_{i,t} = \sigma_\rho^{-2} \text{Tr}(K^\rho) \leq \sigma_\rho^{-2} N$, ($w \log k(v, v) \leq 1$).

From Eq. (49),

$$\begin{aligned} \left(\sum_{i=1}^N \text{Reg}^i(T) \right)^2 &\leq T \sum_{t=1}^T \left(2\sqrt{\beta_t^\rho} \sum_{i=1}^N \sum_{v \in D_t^{i-}} \sigma_{t-1}^\rho(v)/|V| \right)^2 \quad (\text{Cauchy-Schwarz inequality}) \\ &\leq T \sum_{t=1}^T \frac{8C_D N \beta_t^\rho}{\log(1 + N\sigma_\rho^{-2})} \sum_{i=1}^N \frac{1}{2} \log(1 + \sigma_\rho^{-2} \lambda_{i,t}) \quad (\text{Using Eq. (55)}) \\ &= \frac{8C_D N T \beta_T^\rho}{\log(1 + N\sigma_\rho^{-2})} I(Y_{x_{1:T}^{g,1:N}}; \rho) \quad (\text{Since } \beta_t^\rho \text{ is non-decreasing, using Eq. (44)}) \end{aligned}$$

$$\begin{aligned}
&\leq \frac{8C_D N T \beta_T^\rho \gamma_{NT}^\rho}{\log(1 + N\sigma_\rho^{-2})} \quad (\gamma_{NT}^\rho = \sup_{x_{1:T}^{g,1:N} \subset V} I(Y_{x_{1:T}^{g,1:N}}; \rho)) \\
\Rightarrow \sum_{i=1}^N \text{Reg}^i(T) &\leq \sqrt{\frac{8C_D N T \beta_T^\rho \gamma_{NT}^\rho}{\log(1 + N\sigma_\rho^{-2})}} \quad (56)
\end{aligned}$$

Hence Proved. \square

Theorem 1 follows from Lemma 6, Lemma 7 and Eq. (26),

$$\text{Reg}_{act}(T) \leq \sum_{i=1}^N \text{Reg}^i(T) \leq \sqrt{\frac{8C_D N T \beta_T^\rho \gamma_{NT}^\rho}{\log(1 + N\sigma_\rho^{-2})}}$$

Proof for the corollary 1:

Corollary 1. Let t_ρ^* be the smallest integer, such that $\frac{t_\rho^*}{\beta_{t_\rho^*} \gamma_{N t_\rho^*}} \geq \frac{8C_D^2 N^2}{\log(1+N\sigma^{-2})\epsilon_\rho^2}$, then there exists a $t < t_\rho^*$ such that w.h.p, MACOPT terminates and achieves, $F(X_t; \rho, V) \geq (1 - \frac{1}{e})F(X_\star; \rho, V) - \epsilon_\rho$.

Proof. The proof for the corollary goes in the following 2 steps. First, we show that once $w_t \leq \epsilon_\rho$ implies $F(X_t; \rho, V) \geq (1 - \frac{1}{e})F(X_\star; \rho, V) - \epsilon_\rho$. Secondly, in Lemma 8 we show MACOPT achieves $w_t \leq \epsilon_\rho$, at $t < t_\rho^*$ where t_ρ^* be the smallest integer satisfying $\frac{t_\rho^*}{\beta_{t_\rho^*} \gamma_{N t_\rho^*}} \leq \frac{8C_D^2 N^2}{\log(1+N\sigma^{-2})\epsilon_\rho^2}$.

Similar to steps in Lemma 6 for a fix t, (Eq. (48)), we get

$$\begin{aligned}
\Delta(\tilde{x}^i | X_t^{1:i-1}) - \Delta(x_t^i | X_t^{1:i-1}) &\leq 2\sqrt{\beta_t^\rho} \sum_{v \in D_t^{i-}} \sigma_{t-1}^\rho(v) / |V| \\
&\leq 2\sqrt{\beta_t^\rho} C_D \max_{v \in D_t^{i-}} \sigma_{t-1}^\rho(v)
\end{aligned}$$

From Eq. (37) (for constrained case) one can show for unconstrained case,

$$\begin{aligned}
(1 - \frac{1}{e})F(X_\star; \rho, V) - \sum_i \Delta(x_t^i | X_t^{1:i-1}) &\leq \sum_i \Delta(\tilde{x}^i | X_t^{1:i-1}) - \Delta(x_t^i | X_t^{1:i-1}) \\
&\leq \sum_i 2\sqrt{\beta_t^\rho} C_D \max_{v \in D_t^{i-}} \sigma_{t-1}^\rho(v) \leq \epsilon_\rho \\
\Rightarrow F(X_t; \rho, V) &\geq (1 - \frac{1}{e})F(X_\star; \rho, V) - \epsilon_\rho
\end{aligned}$$

\square

Lemma 8. Let $\delta \in (0, 1)$ and β_t^ρ as in [18], i.e., $\beta_t^{\rho^{1/2}} = B_\rho + 4\sigma_\rho \sqrt{\gamma_{Nt}^\rho + 1 + \ln(1/\delta)}$ and t_ρ^* is the smallest integer such that $\frac{t_\rho^*}{\beta_{t_\rho^*} \gamma_{N t_\rho^*}} \geq \frac{8C_D^2 N^2}{\log(1+N\sigma^{-2})\epsilon_\rho^2}$, then with probability $1 - \delta$ that there exists $t_\rho < t_\rho^*$ such that $w_{t_\rho+1} \leq \epsilon_\rho$, where $w_t = 2\sqrt{\beta_t^\rho} C_D \sum_{i=1}^N \max_{v \in D_t^{i-}} \sigma_{t-1}^\rho(v) \leq \epsilon_\rho$.

Proof. It is defined that,

$$\begin{aligned}
w_t &:= 2\sqrt{\beta_t^\rho} C_D \sum_{i=1}^N \max_{v \in D_t^{i-}} \sigma_{t-1}^\rho(v) \\
\Rightarrow w_t^2 &\leq 4\beta_t^\rho C_D^2 N \sum_{i=1}^N \left(\sigma_{t-1}^\rho(x_t^{g,i}) \right)^2 \quad (57)
\end{aligned}$$

$$\Rightarrow \left(\sum_{t=1}^T w_t \right)^2 \leq T \sum_{t=1}^T w_t^2 \leq (C_D N) T \sum_{t=1}^T 4\beta_t^\rho C_D \sum_{i=1}^N \left(\sigma_{t-1}^\rho(x_t^{g,i}) \right)^2 \quad (58)$$

$$\leq (C_D N) \frac{8C_D N T \beta_T^\rho \gamma_{NT}^\rho}{\log(1 + N\sigma_\rho^{-2})} \quad (59)$$

Eq. (57) follows from Cauchy-Schwarz inequality and $x_t^{g,i} = \arg \max_{v \in D_t^{i-}} \sigma_{\rho_{t-1}}(v)$. The RHS of Eq. (57) resembles Eq. (52) in Lemma 7, with an additional factor of $(C_D N)$. Eq. (58) directly follows from Cauchy-Schwarz inequality and Eq. (57). Following the steps in Lemma 7 will result in Eq. (59).

Since,

$$\begin{aligned} \frac{t_\rho^*}{\beta_{t_\rho^*} \gamma_{N t_\rho^*}^\rho} &\geq \frac{8C_D^2 N^2}{\log(1 + N\sigma_\rho^{-2}) \epsilon_\rho^2} \\ \Rightarrow (C_D N)^{1/2} \sqrt{\frac{8C_D N \beta_{t_\rho^*} \gamma_{N t_\rho^*}^\rho}{t_\rho^* \log(1 + N\sigma_\rho^{-2})}} &\leq \epsilon_\rho \quad (\text{Rearranging terms}) \\ \frac{\sum_{t=1}^{t_\rho^*} w_t}{t_\rho^*} &\leq (C_D N)^{1/2} \sqrt{\frac{8C_D N \beta_{t_\rho^*} \gamma_{N t_\rho^*}^\rho}{t_\rho^* \log(1 + N\sigma_\rho^{-2})}} \leq \epsilon_\rho \quad (\text{From Eq. (59)}) \\ \Rightarrow \min_{t \in [1, t_\rho^*]} w_t &\leq \epsilon_\rho \quad \left(\frac{t_\rho^* \min_{t \in [1, t_\rho^*]} w_t}{t_\rho^*} \leq \frac{\sum_{t=1}^{t_\rho^*} w_t}{t_\rho^*} \right) \end{aligned}$$

Hence there exists $t_\rho < t_\rho^*$, such that $w_{t_\rho+1} \leq \epsilon_\rho$.

□

D.1 Variants of MACOPT

- *Hallucinated uncertainty sampling*: Let M_t and H_t be the sets of measurements collected by MACOPT and MACOPT-H respectively at time t . In any iteration, $I(Y_{M_t}; \rho) \leq I(Y_\star; \rho)$, where $I(Y_\star; \rho)$ is the maximum information under the disc constraints. Since mutual information is a submodular function [19], it is a typical submodular maximization under partition matroid constraint (disc constraint). The greedy algorithm (Hallucinated strategy) yields a $1/2$ - times the optimal solution [13]. Hence, using $I(Y_{H_t}; \rho) \geq 1/2 I(Y_\star; \rho)$, we get $I(Y_{M_t}; \rho) \leq 2I(Y_{H_t}; \rho)$. Analogous to Lemma 7, for MACOPT-H we obtain,

$$Reg_{act}^H(T) \leq \sqrt{\frac{16C_D N T \beta_T^\rho \gamma_{NT}^\rho}{\log(1 + N\sigma_\rho^{-2})}}$$

The regret bound worsens by two folds to account for the greedy selection in a partition matroid constraint. But practically, it can improve sample efficiency in environments with high coverage to domain (C_D) ratio.

- *Correlated upper bound*: The coverage function is a linear functional of density. For any agent i , $F(\{x^i\}) = \sum_{v \in D^{i-}} \rho(v)$. Since density is correlated, we can construct a tighter upper bound (in contrast to the sum of density UCB) of the coverage function utilizing the covariance of density. Practically, the sampling rule is given by $x_t^i = \arg \max_v \sum_{v \in D_t^{i-}} \mu_{t-1}^\rho(v) + \sqrt{\beta_t'} \sqrt{\sum_{v \in D_t^{i-}} \sigma_{t-1}^2(v) + \sum_{v, v' \in D_t^{i-}} \sigma_{t-1}(v, v')}$. We believe theoretical analysis for constructing confidence bounds for linear functionals of a sample from RKHS can be carried out utilizing ideas from Mutn y and Krause [66].

E Proof. for Theorem 2 (SAFEMAC)

Theorem 2. Let $\delta \in (0, 1)$, $\epsilon_\rho \geq 0$, $\|\rho\|_{k^\rho} \leq B_\rho$, $\beta_t^{\rho^{1/2}} = B_\rho + 4\sigma_\rho \sqrt{\gamma_{Nt}^\rho + 1 + \ln(1/\delta)}$, γ_{Nt}^ρ denote the information capacity associated with the kernel k^ρ . Let $q(\cdot)$ be L_q -Lipschitz continuous and $\epsilon_q, \beta_t^q, \gamma_{Nt}^q$ be defined analogously. Given $X_0 \neq \emptyset$, $q(x_0^i) \geq 0$ for all $i \in [N]$. Then, for any heuristic $h_t : V \rightarrow \mathbb{R}$, with probability at least $1 - \delta$, we have $q(x) \geq 0$, for any x along the state trajectory pursued by any agent in SAFEMAC. Moreover, let t_ρ^* be the smallest integer such that $\frac{t_\rho^*}{\beta_{t_\rho^*} \gamma_{Nt_\rho^*}} \geq \frac{8C_D^2 N^2}{\log(1+N\sigma_\rho^{-2})\epsilon_\rho^2}$, with $C_D = \max_{x^i \in V} \frac{|D^i|}{|V|} \leq 1$ and let t_q^* be the smallest integer such that $\frac{t_q^*}{\beta_{t_q^*} \gamma_{Nt_q^*}} \geq \frac{C|\bar{R}_0(X_0)|}{\epsilon_q^2}$, with $C = 8/\log(1 + \sigma_q^{-2})$ then, there exists $t \leq t_q^* + t_\rho^*$, such that with probability at least $1 - \delta$,

$$\sum_{B \in \mathcal{B}_t} F(X_t^B; \rho, \bar{R}_0(X_0^B)) \geq \left(1 - \frac{1}{e}\right) \sum_{B \in \mathcal{B}} F(X_\star^B; \rho, \bar{R}_{\epsilon_q}(X_0^B)) - \epsilon_\rho. \quad (7)$$

Proof. The proof for Theorem 2 goes in the following two steps:

1. SAFEMAC's coverage is near-optimal at the convergence

- We first bound the actual regret with the sum of per agent regret in Lemma 4. Precisely, we show the following (Eq. (38)),

$$Reg_{act}^O(T_\rho) \leq Reg_l^O(T_\rho)$$

- Next, we establish in Lemma 9 that the $Reg_l^O(T_\rho)$ grows sublinear with the density measurements.
- Next, we show that if $w_t < \epsilon_\rho$, the coverage is near optimal (Lemma 10). The condition $w_t < \epsilon_\rho$ will eventually happen since $Reg_l^O(T_\rho)$ is sublinear and hence over time will shrink to zero.
- Finally using Lemma 14, the near optimality in the pessimistic set can be established at convergence when the 2nd termination condition is satisfied, precisely $\{S_t^{o, \epsilon_q, i} \setminus S_t^{p, i}\} \cap D_t^i, \forall i \in [N]\} = \emptyset$

2. SAFEMAC converges in a finite time $t < t_q^* + t_\rho^*$, where t_ρ^* be the smallest integer such that

$$\frac{t_\rho^*}{\beta_{t_\rho^*} \gamma_{Nt_\rho^*}} \geq \frac{8C_D^2 N^2}{\log(1+N\sigma_\rho^{-2})\epsilon_\rho^2} \text{ and } t_q^* \text{ be the smallest integer such that } \frac{t_q^*}{\beta_{t_q^*} \gamma_{Nt_q^*}} \geq \frac{C_1 |\bar{R}_0(X_0)|}{\epsilon_q^2},$$

with $C_1 = 8/\log(1 + \sigma_q^{-2})$.

- Since SAFEMAC runs by iterating between the coverage and the exploration phase, we decouple it and analyze both the phases separately. Starting with the *coverage phase*, In Proposition 3, we establish a bound on density samples required to terminate the first coverage phase
- Next, in the Lemma 11, we show that cumulative regret grows sublinear with the density measurements for any coverage phase and utilizes this to bound the density samples between two consecutive coverage phases in Lemma 12
- Utilizing the above two statements, we present the sample complexity bound to terminate the n^{th} coverage phase till convergence, using that the information gain is additive for consecutive coverage phases in Lemma 13
- For the *exploration phase*, the worst case time complexity bound is given by the multi-agent version of the GoOSE in Lemma 19 when the agents safely explore the complete domain. The resulting worst case time bound for SAFEMAC is sum of the time bound of the *coverage* and the *exploration* phase.

So, near optimality at convergence in Theorem 2 is a direct consequence of Lemma 10 and Lemma 14 and the finite time argument of Theorem 2 is a direct consequence of Lemma 13 and Lemma 19. \square

Lemma 9. Let $\delta \in (0, 1)$ and β_t^ρ as in [18], i.e., $\beta_t^{\rho^{1/2}} = B_\rho + 4\sigma\sqrt{\gamma_t^\rho + 1 + \ln(1/\delta)}$. With probability at least $1 - \delta$, SAFEMAC's sum of per agent regret $\text{Reg}_l^O(T_\rho)$ is bounded by $\mathcal{O}(\sqrt{T_\rho\beta_T^\rho\gamma_{NT}^\rho})$. Precisely,

$$\text{Reg}_l^O(T_\rho) \leq \sqrt{\frac{8C_D N T_\rho \beta_T^\rho \gamma_{NT}^\rho}{\log(1 + N\sigma_\rho^{-2})}}$$

where T_ρ is density samples per agent and $\text{Reg}_l^O(T_\rho) = \sum_{t=1}^{T_\rho} r_t$ where $r_t = \sum_{B \in \mathcal{B}_t} \sum_{i \in B} \Delta(\tilde{x}^i | X_t^{1:i-1}; \rho, S_t^{u,B}) - \Delta(x_t^i | X_t^{1:i-1}; \rho, S_t^{u,B})$

Proof. **Given.**

$$\begin{aligned} \text{Reg}_l^O(T_\rho) &= \sum_{t=1}^{T_\rho} r_t \\ &= \sum_{t=1}^{T_\rho} \sum_{B \in \mathcal{B}_t} \sum_{i \in B} \Delta(\tilde{x}^i | X_t^{1:i-1}; \rho, S_t^{u,B}) - \Delta(x_t^i | X_t^{1:i-1}; \rho, S_t^{u,B}) \end{aligned}$$

WLOG, every batch B , is indexed by iterator $i = 1$ to $|B|$ sequentially.

Let $\tilde{x}^i = \arg \max \Delta(x_t^i | X_t^{1:i-1}; \rho, S_t^{u,B})$ and \tilde{D}_t^i is a disk around \tilde{x}^i . For notation convenience: $D_t^{i-} := D_t^i \setminus D_t^{1:i-1} \cap S_t^{u,B}$ and $\tilde{D}_t^{i-} := \tilde{D}_t^i \setminus D_t^{1:i-1} \cap S_t^{u,B}$

SAFEMAC picks the agent at x_t^i greedily in the set B . Following the steps in Lemma 10 we can bound simple agent-wise local regret as r_t or simply from Eq. (66) by summing over all the $B \in \mathcal{B}_t$, we get,

$$\begin{aligned} r_t &= \sum_{B \in \mathcal{B}_t} \sum_{i \in B} \Delta(\tilde{x}^i | X_t^{1:i-1}; \rho, S_t^{u,B}) - \Delta(x_t^i | X_t^{1:i-1}; \rho, S_t^{u,B}) \\ &\leq \sum_{B \in \mathcal{B}_t} \sum_{i \in B} 2\sqrt{\beta_t^\rho} \sum_{v \in D_t^{i-}} \sigma_{t-1}^\rho(v) / |V| \quad (\text{From Eq. (65)}) \\ &\leq \sum_{B \in \mathcal{B}_t} \sum_{i \in B} 2\sqrt{\beta_t^\rho} C_D \max_{v \in D_t^{i-}} \sigma_{t-1}^\rho(v) = w_t \quad (\text{From Eq. (66)}) \end{aligned}$$

On bounding simple regret.

$$\begin{aligned} r_t &\leq \sum_{B \in \mathcal{B}_t} \sum_{i \in B} 2\sqrt{\beta_t^\rho} \sum_{v \in D_t^{i-}} \sigma_{t-1}^\rho(v) / |V| \\ \Rightarrow r_t^2 &\leq \left(\sum_{B \in \mathcal{B}_t} \sum_{i \in B} 2\sqrt{\beta_t^\rho} \sum_{v \in D_t^{i-}} \sigma_{t-1}^\rho(v) / |V| \right)^2 \\ &\leq 4\beta_t^\rho \sum_{B \in \mathcal{B}_t} \sum_{i \in B} \sum_{v \in D_t^{i-}} \left(\sigma_{t-1}^\rho(v) \right)^2 / |V| = 4\beta_t^\rho \sum_{i=1}^N \sum_{v \in D_t^{i-}} \left(\sigma_{t-1}^\rho(v) \right)^2 / |V| \quad (60) \\ &\leq \frac{8C_D N \beta_t^\rho}{\log(1 + N\sigma_\rho^{-2})} \sum_{i=1}^N \frac{1}{2} \log(1 + \sigma_\rho^{-2} \lambda_{i,t}) \quad (61) \end{aligned}$$

Eq. (60) follows by Cauchy-Schwarz inequality and $\sum_{B \in \mathcal{B}_t} \sum_{i \in B} |D_t^{i-}| \leq |V|$. Eq. (61) follows the steps in Eqs. (50) to (55).

On bounding cumulative regret with mutual information.

$$\left(\sum_{t=1}^{T_\rho} r_t \right)^2 \leq T_\rho \sum_{t=1}^{T_\rho} r_t^2 \quad (\text{Using Cauchy-Schwarz inequality})$$

$$\begin{aligned}
&\leq T_\rho \sum_{t=1}^{T_\rho} \frac{8C_D N \beta_t^\rho}{\log(1 + N\sigma_\rho^{-2})} \sum_{i=1}^N \frac{1}{2} \log(1 + \sigma_\rho^{-2} \lambda_{i,t}) \quad (\text{Using Eq. (61)}) \\
&= \frac{8C_D N T_\rho \beta_T^\rho}{\log(1 + N\sigma_\rho^{-2})} \sum_{t=1}^{T_\rho} \sum_{i=1}^N \frac{1}{2} \log(1 + \sigma_\rho^{-2} \lambda_{i,t}) \\
&\quad (\text{Since } \beta_t^\rho \text{ is non-decreasing \& } \beta_T^\rho := \beta_{T_\rho}^\rho) \\
&= \frac{8C_D N T_\rho \beta_T^\rho I(Y_{x_{1:T_\rho}^{g,1:N}}; \rho)}{\log(1 + N\sigma_\rho^{-2})} \quad (\text{Using Eq. (44)}) \\
&\leq \frac{8C_D N T_\rho \beta_T^\rho \gamma_{NT}^\rho}{\log(1 + N\sigma_\rho^{-2})} \quad (\gamma_{NT}^\rho = \sup_{X_{1:T_\rho}^m \subset V} I(Y_{X_{1:T_\rho}^m}; \rho)) \\
&\Rightarrow \sum_{t=1}^{T_\rho} r_t \leq \sqrt{\frac{8C_D N T_\rho \beta_T^\rho \gamma_{NT}^\rho}{\log(1 + N\sigma_\rho^{-2})}} \quad (62) \\
&\Rightarrow \text{Reg}_l^0(T_\rho) \leq \sqrt{\frac{8C_D N T_\rho \beta_T^\rho \gamma_{NT}^\rho}{\log(1 + N\sigma_\rho^{-2})}}
\end{aligned}$$

□

This lemma nicely connects the near optimal coverage in the reachable set i.e., $(1 - \frac{1}{e}) \sum_{B \in \mathcal{B}} F(X_\star^B; \rho, \bar{R}_{\epsilon_q}(X_0^B))$, with the coverage in a possibly disjoint optimistic sets. (Note that the only requirement is that the optimistic set needs to always superset $\bar{R}_{\epsilon_q}(X_0)$).

The agents observe the location only if all the agents can reach the max uncertain point under their disk i.e., $2\sqrt{\beta_t^\rho} \max_{v \in D_t^{i-}} \sigma_{t-1}^\rho(v)$. (Accordingly, information gain is defined, and T_ρ above is a counter when all the agents obtain density measurements).

Lemma 10 (SAFEMAC Near-Optimality). *For any $t \geq 1$, if $w_t \leq \epsilon_\rho$ at SAFEMAC's recommendation X_t then with high probability,*

$$\sum_{B \in \mathcal{B}_t} F(X_t^B; \rho, S_t^{u,B}) \geq (1 - \frac{1}{e}) \sum_{B \in \mathcal{B}} F(X_\star^B; \rho, \bar{R}_{\epsilon_q}(X_0^B)) - \epsilon_\rho,$$

where $w_t = \sum_{B \in \mathcal{B}_t} \sum_{i \in B} 2\sqrt{\beta_t^\rho} C_D \max_{v \in D_t^{i-}} \sigma_{t-1}^\rho(v)$.

Proof. Given. SAFEMAC recommends a location for the agent $i \in B$ greedily in the $S_t^{u,B}$ set as per,

$$x_t^i = \arg \max_v \sum_{v \in D_t^{i-}} \mu_{t-1}^\rho(v) + \sqrt{\beta_t^\rho} \sigma_{t-1}^\rho(v) \quad (63)$$

Let $\tilde{x}_t^i = \arg \max \Delta(x_t^i | X_t^{1:i-1}; \rho, S_t^{u,B})$ and $\tilde{D}_t^{i-} := \tilde{D}_t^i \setminus D_t^{1:i-1} \cap S_t^{u,B}$, where \tilde{D}_t^i is a disk around \tilde{x}_t^i . Based on this picking strategy,

$$\begin{aligned}
\sum_{v \in \tilde{D}_t^{i-}} \rho(v) &\leq \sum_{v \in \tilde{D}_t^{i-}} (\mu_{t-1}^\rho(v) + \sqrt{\beta_t^\rho} \sigma_{t-1}^\rho(v)) \quad (\text{Follows due to upper confidence bound}) \\
&\leq \sum_{v \in D_t^{i-}} (\mu_{t-1}^\rho(v) + \sqrt{\beta_t^\rho} \sigma_{t-1}^\rho(v)) \quad (\text{Since, Eq. (63), } x_t^i \text{ is greedily picked}) \\
\sum_{v \in \tilde{D}_t^{i-}} \rho(v) &\leq \sum_{v \in D_t^{i-}} (\mu_{t-1}^\rho(v) + \sqrt{\beta_t^\rho} \sigma_{t-1}^\rho(v)) \quad (64)
\end{aligned}$$

On bounding simple regret. With definition $r_t = \sum_{B \in \mathcal{B}_t} \sum_{i \in B} \Delta(\tilde{x}_t^i | X_t^{1:i-1}; \rho, S_t^{u,B}) - \Delta(x_t^i | X_t^{1:i-1}; \rho, S_t^{u,B})$. Consider,

$$\Delta(\tilde{x}_t^i | X_t^{1:i-1}; \rho, S_t^{u,B}) - \Delta(x_t^i | X_t^{1:i-1}; \rho, S_t^{u,B})$$

$$\begin{aligned}
&= \left(\sum_{v \in \tilde{D}_t^{i-}} \rho(v) - \sum_{v \in D_t^{i-}} \rho(v) \right) / |V| \quad (\text{Note } D_t^{i-} \text{ and } \tilde{D}_t^{i-}) \\
&\leq \left(\sum_{v \in D_t^{i-}} \left(\mu_{t-1}^\rho(v) + \sqrt{\beta_t^\rho} \sigma_{t-1}^\rho(v) \right) - \sum_{v \in D_t^{i-}} \rho(v) \right) / |V| \\
&\quad (\text{Using Eq. (64)}) \\
&\leq \left(\sum_{v \in D_t^{i-}} \left(\mu_{t-1}^\rho(v) + \sqrt{\beta_t^\rho} \sigma_{t-1}^\rho(v) \right) - \left(\mu_{t-1}^\rho(v) - \sqrt{\beta_t^\rho} \sigma_{t-1}^\rho(v) \right) \right) / |V| \\
&\quad (\text{Since, } \rho(v) \geq \mu_{t-1}^\rho(v) - \sqrt{\beta_t^\rho} \sigma_{t-1}^\rho(v) \forall v) \\
&= 2\sqrt{\beta_t^\rho} \sum_{v \in D_t^{i-}} \sigma_{t-1}^\rho(v) / |V| \quad (65) \\
&\leq 2\sqrt{\beta_t^\rho} C_D \max_{v \in D_t^{i-}} \sigma_{t-1}^\rho(v) \quad (66)
\end{aligned}$$

The last inequality follows since $\sum_{v \in D_t^{i-}} \sigma_{t-1}^\rho(v) \leq |D^{i-}| \max_{v \in D_t^{i-}} \sigma_{t-1}^\rho(v)$ and $\frac{|D_t^{i-}|}{|V|} \leq \frac{\max_i |D_t^i|}{|V|} = C_D$. Now,

$$\begin{aligned}
r_t &= \sum_{B \in \mathcal{B}_t} \sum_{i \in B} \Delta(\tilde{x}_t^i | X_t^{1:i-1}; \rho, S_t^{u,B}) - \Delta(x_t^i | X_t^{1:i-1}; \rho, S_t^{u,B}) \\
&\leq \sum_{B \in \mathcal{B}_t} \sum_{i \in B} 2\sqrt{\beta_t^\rho} C_D \max_{v \in D_t^{i-}} \sigma_{t-1}^\rho(v) \quad (\text{from Eq. (66)}) \\
&= w_t \leq \epsilon_\rho
\end{aligned}$$

From Eq. (37), $(1 - \frac{1}{e}) \sum_{B \in \mathcal{B}} F(X_\star^B; \rho, \bar{R}_{\epsilon_q}(X_0^B)) - \sum_{B \in \mathcal{B}_t} F(X_t^B; \rho, S_t^{u,B}) = r_t^{\text{act}} \leq r_t$

$$\Rightarrow \sum_{B \in \mathcal{B}_t} F(X_t^B; \rho, S_t^{u,B}) \geq (1 - \frac{1}{e}) \sum_{B \in \mathcal{B}} F(X_\star^B; \rho, \bar{R}_{\epsilon_q}(X_0^B)) - \epsilon_\rho$$

□

Proposition 3. Let $\delta \in (0, 1)$ and β_t^ρ as in [18], i.e., $\beta_t^{1/2} = B_\rho + 4\sigma_\rho \sqrt{\gamma_{Nt}^\rho + 1 + \ln(1/\delta)}$ and $t_\rho^{\star 1}$ is the smallest integer such that $\frac{t_\rho^{\star 1}}{\beta_{t_\rho^{\star 1}}^\rho I(Y_{x_{g,1:N}^\rho; \rho})} \geq \frac{8C_D^2 N^2}{\log(1+N\sigma_\rho^{-2})\epsilon_\rho^2}$, then with probability $1 - \delta$ that there exists $t_\rho^1 < t_\rho^{\star 1}$ such that $w_{t_\rho^1+1} \leq \epsilon_\rho$, where $w_t = \sum_{B \in \mathcal{B}_t} \sum_{i \in B} 2\sqrt{\beta_t^\rho} C_D \max_{v \in D_t^{i-}} \sigma_{t-1}^\rho(v) \leq \epsilon_\rho$.

Proof. Similar to Eq. (59), It is defined that,

$$\begin{aligned}
w_t &:= 2\sqrt{\beta_t^\rho} C_D \sum_{i=1}^N \max_{v \in D_t^{i-}} \sigma_{t-1}^\rho(v) \\
\Rightarrow w_t^2 &\leq 4\beta_t^\rho C_D^2 N \sum_{i=1}^N \left(\sigma_{t-1}^\rho(x_t^{g,i}) \right)^2 \quad (67)
\end{aligned}$$

$$\Rightarrow \left(\sum_{t=1}^{T_\rho} w_t \right)^2 \leq T_\rho \sum_{t=1}^{T_\rho} w_t^2 \leq (C_D N) T_\rho \sum_{t=1}^{T_\rho} 4\beta_t^\rho C_D \sum_{i=1}^N \left(\sigma_{t-1}^\rho(x_t^{g,i}) \right)^2 \quad (68)$$

$$\leq (C_D N) \frac{8C_D N T_\rho \beta_{T_\rho}^\rho I(Y_{x_{1:T_\rho}^{g,1:N}; \rho})}{\log(1+N\sigma_\rho^{-2})} \quad (69)$$

Eq. (67) follows from Cauchy-Schwarz inequality and $x_t^{g,i} = \arg \max_{v \in D_t^{i-}} \sigma_{\rho_{t-1}}(v)$. The RHS of Eq. (67) resembles Eq. (52) in Lemma 7, with an additional factor of $(C_D N)$. Eq. (68) directly follows from Cauchy-Schwarz inequality and Eq. (67). Following the steps in Lemma 9 will result in Eq. (69).

Since, it is given that

$$\begin{aligned} \frac{t_\rho^{\star 1}}{\beta_{t_\rho^{\star 1}}^\rho I(Y_{x_{1:t_\rho^{\star 1}}^{g,1:N}}; \rho)} &\geq \frac{8C_D^2 N^2}{\log(1 + N\sigma_\rho^{-2})\epsilon_\rho^2} \quad (70) \\ \implies (C_D N)^{1/2} \sqrt{\frac{8C_D N \beta_{t_\rho^{\star 1}}^\rho I(Y_{x_{1:t_\rho^{\star 1}}^{g,1:N}}; \rho)}{t_\rho^{\star 1} \log(1 + N\sigma_\rho^{-2})}} &\leq \epsilon_\rho \quad (\text{Rearranging terms}) \\ \frac{\sum_{t=1}^{t_\rho^{\star 1}} w_t}{t_\rho^{\star 1}} &\leq (C_D N)^{1/2} \sqrt{\frac{8C_D N \beta_{t_\rho^{\star 1}}^\rho I(Y_{x_{1:t_\rho^{\star 1}}^{g,1:N}}; \rho)}{t_\rho^{\star 1} \log(1 + N\sigma_\rho^{-2})}} \leq \epsilon_\rho \quad (\text{From Eq. (62) in Lemma 9}) \\ \implies \min_{t \in [1, t_\rho^{\star 1}]} w_t &\leq \epsilon_\rho \quad \left(\frac{t_\rho^{\star 1} \min_{t \in [1, t_\rho^{\star 1}]} w_t}{t_\rho^{\star 1}} \leq \frac{\sum_{t=1}^{t_\rho^{\star 1}} w_t}{t_\rho^{\star 1}} \right) \end{aligned}$$

Hence there exists $t_\rho^1 < t_\rho^{\star 1}$, such that $w_{t_\rho^1+1} \leq \epsilon_\rho$. \square

For notation convenience we denote with $Reg_l^O(\delta t_\rho^{\star n}) := Reg_l^O(t_\rho^{n-1} + \delta t_\rho^{\star n}) - Reg_l^O(t_\rho^{n-1}) = \sum_{t=t_\rho^{n-1}+1}^{t_\rho^{n-1}+\delta t_\rho^{\star n}} r_t$ and $I(Y_{\delta t_\rho^{\star n}}; \rho) = I(Y_{x_{t_\rho^{n-1}+1:t_\rho^{n-1}+\delta t_\rho^{\star n}}^{g,1:N}}; \rho)$.

Lemma 11. *Let the coverage phase be terminated for the $(n-1)^{th}$ time at t_ρ^{n-1} , and $\delta t_\rho^{\star n}$ be the maximum number of density measurements required to terminate coverage phase for the n^{th} time. Let $\delta \in (0, 1)$ and β_t^ρ as in [18], i.e., $\beta_t^{\rho^{1/2}} = B_\rho + 4\sigma_\rho \sqrt{\gamma_{Nt}^\rho + 1 + \ln(1/\delta)}$, then with probability at least $1 - \delta$ the following inequality holds,*

$$Reg_l^O(\delta t_\rho^{\star n}) \leq \left(\delta t_\rho^{\star n} \sum_{t=t_\rho^{n-1}+1}^{t_\rho^{n-1}+\delta t_\rho^{\star n}} w_t^2 \right)^{1/2} \leq (C_D N)^{1/2} \sqrt{\frac{8\delta t_\rho^{\star n} C_D N \beta_{t_\rho^{n-1}+\delta t_\rho^{\star n}}^\rho I(Y_{\delta t_\rho^{\star n}}; \rho)}{\log(1 + N\sigma_\rho^{-2})}}$$

Proof. With definitions,

$$\begin{aligned} Reg_l^O(\delta t_\rho^{\star n}) &= \sum_{t=t_\rho^{n-1}+1}^{t_\rho^{n-1}+\delta t_\rho^{\star n}} r_t \leq \sum_{t=t_\rho^{n-1}+1}^{t_\rho^{n-1}+\delta t_\rho^{\star n}} w_t \\ \implies (Reg_l^O(\delta t_\rho^{\star n}))^2 &\leq \sum_{t=t_\rho^{n-1}+1}^{t_\rho^{n-1}+\delta t_\rho^{\star n}} w_t \leq \delta t_\rho^{\star n} \sum_{t=t_\rho^{n-1}+1}^{t_\rho^{n-1}+\delta t_\rho^{\star n}} w_t^2 \quad (\text{Using, Cauchy-Schwarz inequality}) \end{aligned}$$

Now, the RHS of the inequality can be simplified as,

$$\begin{aligned} \delta t_\rho^{\star n} \sum_{t=t_\rho^{n-1}+1}^{t_\rho^{n-1}+\delta t_\rho^{\star n}} w_t^2 &\leq (C_D N) \delta t_\rho^{\star n} \sum_{t=t_\rho^{n-1}+1}^{t_\rho^{n-1}+\delta t_\rho^{\star n}} \frac{8C_D N \beta_t^\rho}{\log(1 + N\sigma_\rho^{-2})} \sum_{i=1}^N \frac{1}{2} \log(1 + \sigma_\rho^{-2} \lambda_{i,t}) \quad (\text{using Eq. (68)}) \\ &\leq \frac{8\delta t_\rho^{\star n} N^2 \beta_{t_\rho^{n-1}+\delta t_\rho^{\star n}}^\rho}{\log(1 + N\sigma_\rho^{-2})} \sum_{t=t_\rho^{n-1}+1}^{t_\rho^{n-1}+\delta t_\rho^{\star n}} \sum_{i=1}^N \frac{1}{2} \log(1 + \sigma_\rho^{-2} \lambda_{i,t}) \\ &\quad (\text{since, } \beta_t^\rho \text{ is non-decreasing and using definition of mutual information we get,}) \\ \implies Reg_l^O(\delta t_\rho^{\star n}) &\leq \sum_{t=t_\rho^{n-1}+1}^{t_\rho^{n-1}+\delta t_\rho^{\star n}} w_t \leq \left(\delta t_\rho^{\star n} \sum_{t=t_\rho^{n-1}+1}^{t_\rho^{n-1}+\delta t_\rho^{\star n}} w_t^2 \right)^{1/2} \leq (C_D N)^{1/2} \sqrt{\frac{8\delta t_\rho^{\star n} C_D N \beta_{t_\rho^{n-1}+\delta t_\rho^{\star n}}^\rho I(Y_{\delta t_\rho^{\star n}}; \rho)}{\log(1 + N\sigma_\rho^{-2})}} \quad (71) \end{aligned}$$

□

Lemma 12. Let $\delta \in (0, 1)$ and β_t^ρ as in [18], i.e., $\beta_t^{\rho^{1/2}} = B_\rho + 4\sigma_\rho \sqrt{\gamma_{Nt}^\rho + 1 + \ln(1/\delta)}$ and δt_ρ^{*n} is the smallest integer after t_ρ^{n-1} such that $\frac{\delta t_\rho^{*n}}{\beta_{t_\rho^{n-1} + \delta t_\rho^{*n}}^\rho I(Y_{\delta t_\rho^{*n}}; \rho)} \geq \frac{8C_D^2 N^2}{\log(1 + N\sigma_\rho^{-2})\epsilon_\rho^2}$, then we know with probability $1 - \delta$ that there exists $\delta t_\rho^n < \delta t_\rho^{*n}$ such that $w_{t_\rho^{n-1} + \delta t_\rho^n + 1} \leq \epsilon_\rho$, where $w_t = \sum_{B \in \mathcal{B}_t} \sum_{i \in B} 2\sqrt{\beta_t^\rho} C_D \max_{v \in D_t^i} \sigma_{t-1}^\rho(v) \leq \epsilon_\rho$.

Proof. Given,

$$\begin{aligned} & \frac{\delta t_\rho^{*n}}{\beta_{t_\rho^{n-1} + \delta t_\rho^{*n}}^\rho I(Y_{\delta t_\rho^{*n}}; \rho)} \geq \frac{8C_D^2 N^2}{\log(1 + N\sigma_\rho^{-2})\epsilon_\rho^2} \\ \Rightarrow & (C_D N)^{1/2} \sqrt{\frac{8C_D N \beta_{t_\rho^{n-1} + \delta t_\rho^{*n}}^\rho I(Y_{\delta t_\rho^{*n}}; \rho)}{\delta t_\rho^{*n} \log(1 + N\sigma_\rho^{-2})}} \leq \epsilon_\rho \\ & \frac{\sum_{t_\rho^{n-1} + 1}^{\delta t_\rho^{*n}} w_t}{\delta t_\rho^{*n}} \leq \epsilon_\rho \quad (\text{Using Eq. (71) in Lemma 11}) \\ \Rightarrow & \min_{t \in [t_\rho^{n-1} + 1, t_\rho^{n-1} + \delta t_\rho^{*n}]} w_t \leq \epsilon_\rho \end{aligned}$$

Hence there exists $\delta t_\rho^n < \delta t_\rho^{*n}$, such that $w_{t_\rho^{n-1} + \delta t_\rho^n + 1} \leq \epsilon_\rho$. □

Lemma 13. Let $\delta \in (0, 1)$ and $\beta_t^{\rho^{1/2}} = B_\rho + 4\sigma_\rho \sqrt{\gamma_{Nt}^\rho + 1 + \ln(1/\delta)}$ and t_ρ^* is the smallest integer such that $\frac{t_\rho^*}{\beta_{t_\rho^*}^\rho \gamma_{Nt_\rho^*}^\rho} \geq \frac{8C_D^2 N^2}{\log(1 + N\sigma_\rho^{-2})\epsilon_\rho^2}$, then for any $n \geq 1$, $t_\rho^{n-1} + \delta t_\rho^n < t_\rho^*$.

Proof.

$$\begin{aligned} t_\rho^{n-1} + \delta t_\rho^n & < \frac{8C_D^2 N^2 \beta_{t_\rho^{n-1}}^\rho I(Y_{x_{1:t_\rho^{n-1}}^{g,1:N}}; \rho)}{\log(1 + N\sigma_\rho^{-2})\epsilon_\rho^2} + \frac{8C_D^2 N^2 \beta_{t_\rho^{n-1} + \delta t_\rho^n}^\rho I(Y_{\delta t_\rho^n}; \rho)}{\log(1 + N\sigma_\rho^{-2})\epsilon_\rho^2} \\ & \quad (\text{using Eq. (70), since } t_\rho^1 < t_\rho^{*1}) \\ & < \frac{8C_D^2 N^2 \beta_{t_\rho^{n-1} + \delta t_\rho^n}^\rho}{\log(1 + N\sigma_\rho^{-2})\epsilon_\rho^2} (I(Y_{x_{1:t_\rho^{n-1}}^{g,1:N}}; \rho) + I(Y_{\delta t_\rho^n}; \rho)) \\ & \quad (\text{Since, } \beta_t^\rho \text{ is non decreasing function}) \\ & = \frac{8C_D^2 N^2 \beta_{t_\rho^{n-1} + \delta t_\rho^n}^\rho I(Y_{x_{1:t_\rho^{n-1} + \delta t_\rho^n}^{g,1:N}}; \rho)}{\log(1 + N\sigma_\rho^{-2})\epsilon_\rho^2} \quad (\text{Since mutual info is additive}) \\ & < \frac{8C_D^2 N^2 \beta_{t_\rho^{n-1} + \delta t_\rho^n}^\rho \gamma_{N(t_\rho^{n-1} + \delta t_\rho^n)}^\rho}{\log(1 + N\sigma_\rho^{-2})\epsilon_\rho^2} \quad (72) \end{aligned}$$

Using Eq. (72) and since, $t_\rho^* \geq \frac{8C_D^2 N^2 \beta_{t_\rho^*}^\rho \gamma_{Nt_\rho^*}^\rho}{\log(1 + N\sigma_\rho^{-2})\epsilon_\rho^2}$, we get $t_\rho^{n-1} + \delta t_\rho^n < t_\rho^*$. □

Lemma 14. When SAFEMAC converges, i.e, $U := \{S_t^{o, \epsilon_q, i} \setminus S_t^{p, i} \cap D_t^i, \forall i \in [N]\} = \emptyset$, then the following inequality holds,

$$\sum_{B \in \mathcal{B}_t} F(X_t^B; \rho, S_t^{p, B}) = \sum_{B \in \mathcal{B}_t} F(X_t^B; \rho, S_t^{u, B})$$

Proof. Since, $U = \emptyset$,

$$\{(S_t^{o, \epsilon_q, i} \setminus S_t^{p, i}) \cap D_t^i, \forall i \in [N]\} = \emptyset$$

$$\begin{aligned}
\Rightarrow (S_t^{o, \epsilon_q, i} \cap D_t^i) &\subseteq (S_t^{p, i} \cap D_t^i) \quad \forall i \in [N] \\
&= (S_t^{u, i} \cap D_t^i) \quad \forall i \in [N] \quad (\text{Since } S_t^{u, i} := S_t^{p, i} \cup S_t^{o, \epsilon_q, i})
\end{aligned}$$

Based on the last equality, it directly follows,

$$\sum_{B \in \mathcal{B}_t} F(X_t^B; \rho, S_t^{p, B}) = \sum_{B \in \mathcal{B}_t} F(X_t^B; \rho, S_t^{u, B}).$$

□

E.1 Intermediate recommendation is near-optimal at SAFEMAC's convergence

Lemma 15. Let $\delta \in (0, 1)$ and β_t^ρ as in [18], i.e., $\beta_t^{\rho^{1/2}} = B_\rho + 4\sigma_\rho \sqrt{\gamma_{N_t}^\rho + 1 + \ln(1/\delta)}$ and t_ρ^* be the smallest integer such that $\frac{t_\rho^*}{\beta_{t_\rho^*}^\rho \gamma_{N_{t_\rho^*}}^\rho} \geq \frac{8C_D^2 N^2}{\log(1+N\sigma_\rho^{-2})\epsilon_\rho^2}$. Let β_t^q and t_q^* be defined analogously. Then, there exists $t < t_q^* + t_\rho^*$, such that with probability at least $1 - \delta$

$$\sum_{B \in \mathcal{B}_T} F(X_T^B; \rho, \bar{R}_0(X_0^B)) \geq (1 - \frac{1}{e}) \sum_{B \in \mathcal{B}} F(X_\star^B; \rho, \bar{R}_{\epsilon_q}(X_0^B)) - \epsilon_\rho \quad (73)$$

where,

$$X_T = \arg \max_{X_T, X_T^l, T \in [1, t]} \left\{ \sum_{B \in \mathcal{B}_T^p} F(X_T^B; l_{T-1}^\rho, S_T^{p, B}), \sum_{B \in \mathcal{B}_T^p} F(X_T^{l, B}; l_{T-1}^\rho, S_T^{p, B}) \right\} \text{ s.t. } X_T \in S_T^p \quad (74)$$

and $X_t^{l, B}$, i.e., the greedy solution w.r.t. the worst-case objective, $F(\cdot; l_{t-1}^\rho, S_t^{p, B}) \forall B \in \mathcal{B}_t^p$.

Proof. We prove the lemma in two parts. First, we prove the near optimality of SAFEMAC's solution X_t but evaluated using l_{t-1}^ρ instead of ρ . This will imply the near optimality at convergence of the 1st term ($\sum_{B \in \mathcal{B}_T^p} F(X_T^B; l_{T-1}^\rho, S_T^{p, B})$) in the above recommendation rule. Secondly, due to the $\arg \max$ operator, the near optimality of the 1st term is sufficient to establish the optimality of the recommendation rule in Eq. (74).

Notations. $X_t = \cup_{B \in \mathcal{B}_t} X_t^B$, $\Delta(\cdot; \rho, V)$ as defined in Eq. (17).

Given. From Theorem 2, for $t < t_q^* + t_\rho^*$ with probability at least $1 - \delta$,

$$\sum_{B \in \mathcal{B}_t} F(X_t^B; \rho, \bar{R}_0(X_0^B)) \geq (1 - \frac{1}{e}) \sum_{B \in \mathcal{B}} F(X_\star^B; \rho, \bar{R}_{\epsilon_q}(X_0^B)) - \epsilon_\rho \quad (75)$$

and

$$\sum_{B \in \mathcal{B}_t} \sum_{i \in B} 2\sqrt{\beta_t^\rho} C_D \max_{v \in D_t^{i-}} \sigma_{t-1}^\rho(v) \leq \epsilon_\rho$$

Near-optimality of SAFEMAC's X_t evaluated using l_{t-1}^ρ .

$$\begin{aligned}
&\Delta(\tilde{x}_t^i | X_t^{1:i-1}; \rho, S_t^{u, B}) - \Delta(x_t^i | X_t^{1:i-1}; l_{t-1}^\rho, S_t^{u, B}) \\
&= \left(\sum_{v \in \tilde{D}_t^{i-}} \rho(v) - \sum_{v \in D_t^{i-}} l_{t-1}^\rho(v) \right) / |V| \quad (\text{Note } D_t^{i-} \text{ and } \tilde{D}_t^{i-}) \\
&\leq \left(\sum_{v \in D_t^{i-}} (\mu_{t-1}^\rho(v) + \sqrt{\beta_t^\rho} \sigma_{t-1}^\rho(v)) - (\mu_{t-1}^\rho(v) - \sqrt{\beta_t^\rho} \sigma_{t-1}^\rho(v)) \right) / |V| \\
&\quad (\text{Using Eq. (64) and definition of } l_{t-1}^\rho(v)) \\
&= 2\sqrt{\beta_t^\rho} \sum_{v \in D_t^{i-}} \sigma_{t-1}^\rho(v) / |V|
\end{aligned}$$

$$\leq 2\sqrt{\beta_t^\rho} C_D \max_{v \in D_t^{i-}} \sigma_{t-1}^\rho(v) \quad (76)$$

$$\begin{aligned} \sum_{B \in \mathcal{B}_t} \sum_{i \in B} \Delta(\tilde{x}_t^i | X_t^{1:i-1}; \rho, S_t^{u,B}) - \Delta(x_t^i | X_t^{1:i-1}; l_{t-1}^\rho, S_t^{u,B}) &\leq \sum_{B \in \mathcal{B}_t} \sum_{i \in B} 2\sqrt{\beta_t^\rho} C_D \max_{v \in D_t^{i-}} \sigma_{t-1}^\rho(v) \\ &\leq \epsilon_\rho \end{aligned}$$

Using the following two statements,

- $(1 - \frac{1}{e})F(X_\star; \rho, S_t^{u,B}) \leq \sum_{i \in B} \Delta(\tilde{x}_t^i | X_t^{1:i-1}; \rho, S_t^{u,B})$ from Eq. (39)
- $\bigcup_{i \in B} \bar{R}_{\epsilon_q}(\{x_0^i\}) \subseteq S_t^{u,B} \implies \sum_{B \in \mathcal{B}} F(X_\star; \rho, \bar{R}_{\epsilon_q}(X_0)) \leq \sum_{B \in \mathcal{B}_t} F(X_\star; \rho, S_t^{u,B})$

we get,

$$\implies \sum_{B \in \mathcal{B}_t} F(X_t^B; l_{t-1}^\rho, S_t^{u,B}) \geq (1 - \frac{1}{e}) \sum_{B \in \mathcal{B}} F(X_\star^B; \rho, \bar{R}_{\epsilon_q}(X_0^B)) - \epsilon_\rho \quad (77)$$

Near-optimality of recommendation as per Eq. (74).

Let's consider the following recommendation rule,

$$X_T = \arg \max_{X_T, T \in [1, t]} \left\{ \sum_{B \in \mathcal{B}_T^p} F(X_T^B; l_{T-1}^\rho, S_T^{p,B}) \right\} \text{ s.t. } X_T \in S_T^p \quad (78)$$

At convergence, $S_t^{p,i} \cap D_t^i = S_t^{u,i} \cap D_t^i \implies (S^{o, \epsilon_q} \setminus S_t^{p,i}) \cap D_t^i = \emptyset$, using this SAFEMAC recommendation X_t can be written as,

$$\begin{aligned} \sum_{B \in \mathcal{B}_t} F(X_t^B; l_{t-1}^\rho, S_t^{u,B}) &= \sum_{i \in [N]} \Delta(x_t^i | X_t^{1:i-1}; l_{t-1}^\rho, S_t^{p,i}) = \sum_{B \in \mathcal{B}_t^p} F(X_t^B; l_{t-1}^\rho, S_t^{p,B}) \\ \sum_{B \in \mathcal{B}_T^p} F(X_T^B; l_{T-1}^\rho, S_T^{p,B}) &\geq \sum_{B \in \mathcal{B}_t^p} F(X_t^B; l_{t-1}^\rho, S_t^{p,B}) \\ &\quad (\text{since, } X_T^B = \arg \max_{X_T^B, T \in [1, t]} \sum_{B \in \mathcal{B}_T^p} F(X_T^B; l_{T-1}^\rho, S_T^{p,B})) \\ \implies \sum_{B \in \mathcal{B}_T^p} F(X_T^B; l_{T-1}^\rho, S_T^{p,B}) &\geq \sum_{B \in \mathcal{B}_t} F(X_t^B; l_{t-1}^\rho, S_t^{u,B}) \quad (\text{Combining the above 2 equations}) \\ \implies \sum_{B \in \mathcal{B}_T^p} F(X_T^B; l_{T-1}^\rho, S_T^{p,B}) &\geq (1 - \frac{1}{e}) \sum_{B \in \mathcal{B}} F(X_\star^B; \rho, \bar{R}_{\epsilon_q}(X_0^B)) - \epsilon_\rho \quad (\text{using Eq. (77)}) \end{aligned}$$

Hence, the recommendation of Eq. (78) evaluated with lower bound is near optimal (at convergence $X_T \in S_T^p$). Further, due to arg max operator Eq. (78) also implies near-optimality of recommendation rule in Eq. (74) evaluated with the lower bound. So now using X_T chosen as per Eq. (74) and at convergence, $\forall i, (S_t^{p,i} \cap D_t^i) \subseteq (\bar{R}_0(\{x_0^i\}) \cap D_t^i)$, we get,

$$\begin{aligned} \sum_{B \in \mathcal{B}_T^p} F(X_T^B; \rho, S_T^{p,B}) &\geq (1 - \frac{1}{e}) \sum_{B \in \mathcal{B}} F(X_\star^B; \rho, \bar{R}_{\epsilon_q}(X_0^B)) - \epsilon_\rho \quad (l_{t-1}^\rho(v) \leq \rho(v) \forall v) \\ \sum_{B \in \mathcal{B}_T^p} F(X_T^B; \rho, \bar{R}_0(X_0)) &\geq (1 - \frac{1}{e}) \sum_{B \in \mathcal{B}} F(X_\star^B; \rho, \bar{R}_{\epsilon_q}(X_0^B)) - \epsilon_\rho \end{aligned}$$

Hence Proved. □

F Multi-agent GoOSE version

In this section, we present our lemma for the multi-agent version of goose. In the cooperative setting, each agent deploy GoOSE for safe exploration and shares its observations with the other agents. We first derive a sample complexity bound under the cooperative system. Later, we introduce our key Lemma 19, which guarantees the safety of all agents as well as complete exploration (with respect to each agent) in finite time.

Lemma 16. *Let $\delta \in (0, 1)$ and let $(\beta_t^q)^{1/2} = B_q + 4\sigma_q \sqrt{\gamma_{Nt}^q + 1 + \ln(1/\delta)}$. Then the following holds with probability at least $1 - \delta$,*

$$\sum_t \omega_t^2 \leq C_1 \beta_t^q I(Y_{NT}; q) \leq C_1 \beta_t^q \gamma_{NT}^q,$$

where $C_1 = 8/\log(1 + \sigma_q^{-2})$, $\omega_t = u_{t-1}^q(x_t^i) - l_{t-1}^q(x_t^i)$, and x_t^i is the location visited by some agent i at time t . $I(Y_{NT}; q)$ is the information gain and γ_{NT}^q is the information capacity.

Proof. Using $\omega_t \leq 2\sqrt{\beta_t^q \sigma_{t-1}^q(x_t^i)}$,

$$\omega_t^2 \leq 4\beta_t^q (\sigma_{t-1}^q(v))^2 \leq 4\beta_t^q \sigma_q^2 \sigma_q^{-2} (\sigma_{t-1}^q(x_t^i))^2 \leq 4\beta_t^q \sigma_q^2 C_2 \log(1 + \sigma_q^{-2} (\sigma_{t-1}^q(x_t^i))^2) \quad (79)$$

$$\leq C_1 \beta_t^q \frac{1}{2} \log(1 + \sigma_q^{-2} \sum_{i=1}^N (\sigma_{t-1}^q(x_t^i))^2) = C_1 \beta_t^q \frac{1}{2} \log(1 + \sigma_q^{-2} \sum_i \lambda_{i,t}) \quad (80)$$

$$\leq C_1 \beta_t^q \sum_{i=1}^N \frac{1}{2} \log(1 + \sigma_q^{-2} \lambda_{i,t}) = C_1 \beta_t^q I(Y_{NT}; q) \leq C_1 \beta_t^q \gamma_{NT}^q \quad (81)$$

Last inequality in Eq. (79) follows since, $s \leq C_2 \log(1 + s)$ for $s \in [0, \sigma_q^{-2}]$, where $C_2 = \sigma_q^{-2} / \log(1 + \sigma_q^{-2}) \geq 1$, where $s = \sigma_q^{-2} \sigma_{t-1}^q(v)^2 \leq \sigma_q^{-2} k^q(v, v) \leq \sigma_q^{-2}$, ($w \log k^q(v, v) \leq 1$). Inequality of Eq. (80) follows since, $(\sigma_{t-1}^q(x_t^i))^2 \leq \sum_{i=1}^N (\sigma_{t-1}^q(x_t^i))^2$ and equality using $\sum_{i=1}^N (\sigma_{t-1}^q(x_t^i))^2 = \text{Tr}(K^q) = \sum_{i=1}^N \lambda_{i,t}$. Eq. (81) follows since $\log(1 + x_1 + x_2) \leq \log(1 + x_1) + \log(1 + x_2)$, for $x_1, x_2 \geq 0$. Lastly, $I(\cdot; q)$ is defined analogous to $I(\cdot; \rho)$ (as in Eq. (44)) and $\gamma_{NT}^q = \sup_{A \subset V; |A|=NT} I(Y_A; q)$. \square

Similar to Lem. 8 of Turchetta et al. [12], Let us denote $\mathcal{T}_t^v = \{\tau_1, \dots, \tau_j\}$ the set of steps where the constraint q is evaluated at v by step t .

Lemma 17. *For any $t \geq 1$ and for any $v \in V$, it holds that $w_t(v) \leq \sqrt{\frac{C_1 \beta_t^q \gamma_{Nt}^q}{|\mathcal{T}_t^v|}}$, with $C_1 = 8/\log(1 + \sigma_q^{-2})$.*

Proof.

$$\begin{aligned} |\mathcal{T}_t^v| w_t^2(v) &\leq \sum_{\tau \in \mathcal{T}_t^v} w_\tau^2(v) \\ &\leq \sum_{\tau \in \mathcal{T}_t^v} 4\beta_\tau^q (\sigma_{t-1}^q(x_t^i))^2 \\ &\leq \sum_{\tau \in \mathcal{T}_t^v} 4\beta_\tau^q (\sigma_{t-1}^q(x_t^i))^2 \\ &\leq C_1 \beta_t^q \gamma_{Nt}^q \end{aligned} \quad (82)$$

Eq. (82), follows due to intersection of confidence interval arguments, Lemma 1 of Turchetta et al. [12] and the inequality follows due to Lemma 16. \square

Let us denote with T_t , the smallest positive integer such that $\frac{T_t}{\beta_{t+T_t}^q \gamma_{N,t+T_t}^q} \geq \frac{C_1}{\epsilon_q^2}$, with $C_1 = 8/\log(1 + \sigma_q^{-2})$ and with t^* the smallest positive integer such that $t^* \geq |\bar{R}_0(X_0)|T_{t^*}$.

Lemma 18. *For any $t \leq t^*$, for any $v \in V$ such that $|\mathcal{T}_t^v| \geq T_{t^*}$, it holds that $w_t(v) \leq \epsilon_q$.*

Proof. Since T_t is an increasing function of t [26], we have $|\mathcal{T}_t^v| \geq T_{t^*} \geq T_t$. Therefore using Lemma 17 and the definition of T_t , we get,

$$w_t(v) \leq \sqrt{\frac{C_1 \beta_t^q \gamma_{Nt}^q}{T_t}} \leq \sqrt{\frac{C_1 \beta_t^q \gamma_{Nt}^q \epsilon_q^2}{C_1 \gamma_{N,t+T_t}^q \beta_{t+T_t}^q}} \leq \sqrt{\frac{\beta_t^q \gamma_{Nt}^q}{\gamma_{N,t+T_t}^q \beta_{t+T_t}^q}} \epsilon_q \leq \epsilon_q.$$

The last inequality follows from the fact that both β_t^q and γ_t^q are non-decreasing function of t . \square

Regarding the convergence of the pessimistic and the optimistic sets, Lemma 10-18 of Turchetta et al. [12] can be proved analogously for each agent i . We skip re-writing them and directly cite them in the following lemma.

Lemma 19. *Assume that $q(\cdot)$ is L_q -Lipschitz continuous w.r.t $d(\cdot, \cdot)$ with $\|q\|_k \leq B_q$, $X_0 \neq \emptyset$, $q(x_0^i) \geq 0$ for all $i \in [N]$. Let $(\beta_t^q)^{1/2} = B_q + 4\sigma_q \sqrt{\gamma_{Nt}^q + 1 + \ln(1/\delta)}$, then, for any heuristic $h_t : V \rightarrow \mathbb{R}$, with probability at least $1 - \delta$, we have $q(x) \geq 0$, for any x along the state trajectory pursued by any agent in SAFEMAC. Moreover, let γ_{Nt}^q denote the information capacity associated with the kernel k^q and let t_q^* be the smallest integer such that $\frac{t_q^*}{\beta_{t_q^*} \gamma_{Nt_q^*}^q} \geq \frac{C_1 |\bar{R}_0(X_0)|}{\epsilon_q^2}$, with $C_1 = 8/\log(1 + \sigma_q^{-2})$, then there exists $t \leq t_q^*$ such that, with probability at least $1 - \delta$, $\bar{R}_{\epsilon_q}(\{x_0^i\}) \subseteq S_t^{o, \epsilon_q, i} \subseteq S_t^{p, i} \subseteq \bar{R}_0(\{x_0^i\})$ for all $i \in [N]$.*

Proof. In SAFEMAC, each agent has a record of its optimistic and pessimistic set. The lemma is similar to N instances of Theorem 1 of Turchetta et al. [12]; each instance corresponds to per agent case. Safety of each agent i is a direct consequence of Theorem 2 of Turchetta et al. [64]. Finite time bound while agents are sharing information is consequence of Lemma 16-18. The convergence of the pessimistic and optimistic approximation of the safe sets for each agent is a direct consequence of Lemmas 16-18 of Turchetta et al. [12]. \square

For a detailed discussion, we refer the reader to Appendix D Completeness of Turchetta et al. [12].

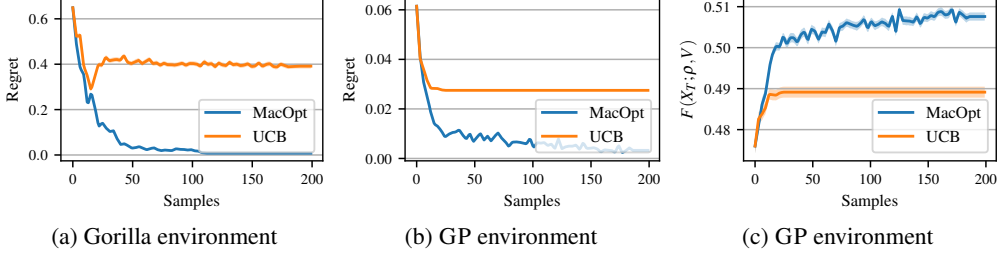


Figure 6: Compares MACOPT and UCB on the Gorilla (a) and the GP (b,c) environment. a,b) Compares simple regret r_t Eq. (29) in the unconstrained case (domain V). c) Plots total coverage achieved by both the algorithms.

G Experiments

Implementation details. We implemented all our algorithms with BoTorch [67] and GPyTorch [68] frameworks, built on top of Pytorch [69]. The code for both the algorithms will be made public along with the competitive baselines. We limit the maximum number of rounds to 300, and with the selected hyperparameters and the given environments, it terminates before that. This roughly takes *10 min* of training for SAFEMAC on a single core CPU. The code is written for running a single instance of the experiment. In practice, we launch nearly 1000 such instances simultaneously on the cluster in parallel to get results about different environments, noise realizations and initializations.

Gorilla Environment. The gorilla environment (Fig. 3b) is defined in a grid of 34×34 , with each grid cell being a square of length 0.1. The $N = 3$ agents perform the coverage task, with each having a sensing region defined as a set of locations agents that can travel in 5 steps in the underlying transition graph (Precisely, $D^i = R_5^{\text{reach}}(\{x^i\})$, Eq. (9)). We considered 10 gorilla environments each differ in the initial location of the agents. The nest density is obtained by fitting a smooth rate function [23] over Gorilla nest site locations which were provided by the Wildlife Conservation Society Takamanda-Mone Landscape project (WCS-TMLP) Funwi-gabga and Mateu [24]. As a proxy for bad weather, we use the cloud coverage data over the Kagwene Gorilla Sanctuary from OpenWeather [22]. The density and the constraint function used are available in our code base. The code for fitting a rate function is available here (<https://github.com/Mojusko/sensepy>) under the MIT license. We used a lengthscale of 1 for the density and of 2 for the constraint function. The noise variance is set to 10^{-3} and 7×10^{-3} for density and the constraint respectively. However, the performance in the experiments is not sensitive to the hyperparameters and is easily reproducible with other sensible parameters as well.

Obstacles Environment. The obstacle environment (Fig. 3a) is defined on a grid of 30×30 , with each grid cell being a square of length 0.1. The sensing region and number of agents are defined similar to the Gorilla environment. The obstacle is completely defined by the location of its top right corner and the bottom left corner. The obstacle environment is generated by combining a set of such obstacles. The density is directly sampled from a GP with the parameters same as synthetic data. We produced ten instances of environments, each having a different set of obstacles and GP sample and initialization. We used a lengthscale of 2 for both density and the constraint function. The noise variance is set to 10^{-3} . Similar to earlier environments, performance is not sensitive to hyperparameters.

Experiment results.

Unconstrained case Fig. 6a and Fig. 6b plots the simple regret r_t for each round t , precisely, defined as $\sum_{i=1}^N \Delta(\tilde{x}|X^{1:i-1}; \rho, V) - \Delta(x_t^i|X^{1:i-1}; \rho, V)$. This quantity upper bounds the actual regret and provides intuition for the convergence rate. We see in the plots that the simple regret goes to zero for MACOPT, but gets stuck for the UCB algorithm. Due to this, we also observe that MACOPT can achieve higher coverage value as compared to UCB in Fig. 6c.

Constrained case Fig. 7a and Fig. 7b compares coverage of area attained by SAFEMAC, PASSIVEMAC and the two stage algorithm. Precisely the intermediate locations are recommended as per Eq. (74). We see that SAFEMAC finds a comparable solution to two stage more efficiently without exploring the whole environment, where as PASSIVEMAC gets stuck in the local optimum.

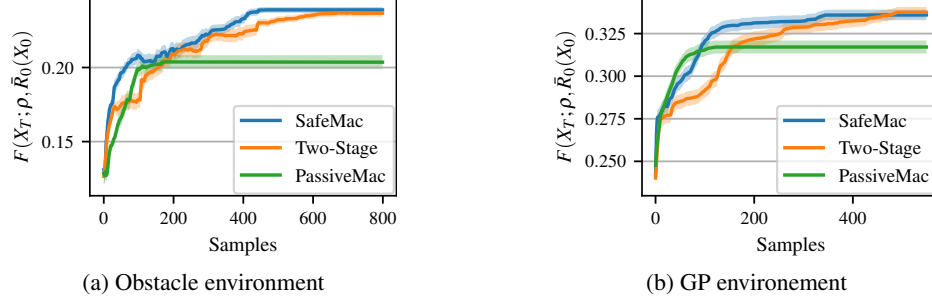


Figure 7: Comparison of SAFEMAC with PASSIVEMAC and Two-Stage in Obstacle and the GP environment during optimization

G.1 Scaling in terms of agents and domain size

In this section, we evaluate scalability in terms of number of agents and the domain size. SAFEMAC evaluates a greedy solution N times (one for each agent) at each iteration, it is linear in the number of agents. Moreover, the greedy solution is linear in the number of cells (domain size). To demonstrate this we run the experiment on the Gorilla nest density for $N = 3, 6, 10$ and 15 agents each for the domains of size $30 \times 30, 40 \times 40, 50 \times 50$ and 60×60 . We see that with more agents in larger domain the same results hold that is SAFEMAC finds a comparable solution to two stage more efficiently without exploring the whole environment, where as PASSIVEMAC gets stuck in the local optimum.

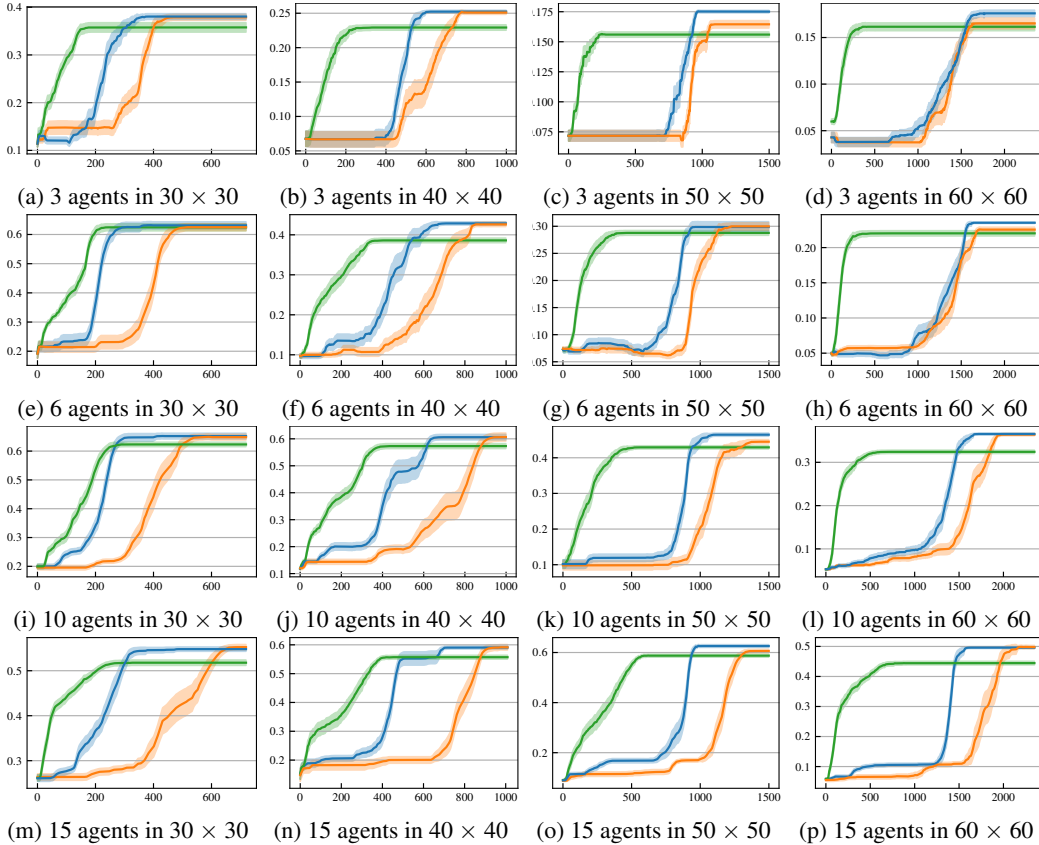


Figure 8: Comparison of SAFEMAC with PASSIVEMAC and Two-Stage in the Gorilla nest environment. Legend: Blue is SAFEMAC, Green is PASSIVEMAC and Orange is Two-Stage algorithm. The experiment is performed for 3, 6, 10 and 15 agents (increased row-wise) each for the domains of size $30 \times 30, 40 \times 40, 50 \times 50$ and 60×60 (increased column-wise).