



Scalable multirobot planning for informed spatial sampling

Sandeep Manjanna¹ · M. Ani Hsieh¹ · Gregory Dudek²

Received: 3 September 2021 / Accepted: 17 May 2022 / Published online: 24 August 2022
© The Author(s), under exclusive licence to Springer Science+Business Media, LLC, part of Springer Nature 2022

Abstract

This paper presents a distributed scalable multi-robot planning algorithm for informed sampling of quasistatic spatial fields. We address the problem of efficient data collection using multiple autonomous vehicles and consider the effects of communication between multiple robots, acting independently, on the overall sampling performance of the team. We focus on the distributed sampling problem where the robots operate independent of their teammates, but have the ability to communicate their current state to other neighbors within a fixed communication range. Our proposed approach is scalable and adaptive to various environmental scenarios, changing robot team configurations, and runs in real-time, which are important features for many real-world applications. We compare the performance of our proposed algorithm to baseline strategies through simulated experiments that utilize models derived from both synthetic and field deployment data. The results show that our sampling algorithm is efficient even when robots in the team are operating with a limited communication range, thus demonstrating the scalability of our method in sampling large-scale environments.

Keywords Environment monitoring · Adaptive sampling · Multi-Robot systems · Marine robots

1 Introduction

In this paper, we present a distributed planning approach to design paths for multiple robots to achieve efficient sampling of a quasistatic spatial field. Our objective is to plan a non-myopic path for the robots to maximize the information gain in a limited time. We address the problem of efficient sampling of environmental processes using multiple autonomous vehicles.

Many natural processes can be modeled by hotspots exhibiting extreme measurements and higher spatial vari-

ability than the rest of the field. Examples of such spatial fields include algal blooms, coral reefs (see Fig. 1), distribution of flora and fauna, and aerosol concentration. To model such environmental phenomena with high precision, we need the sampling robots to visit many information rich locations within their limited endurance. One possible way to achieve this is by exploring the hotspot regions in the early stages of the survey as the hotspots are rich with the information needed to model the phenomenon of interest. Estimating a good representation of spatial phenomena plays a key role in applications like environmental monitoring, search and rescue, anomaly detection, and geological surveys. We propose an informed non-myopic path planning technique for robotic platforms to efficiently collect measurements from a quasistatic spatial field, so that an accurate model of the underlying physical phenomenon can be built.

In a multi-robot setup, task division between the robots becomes essential as opposed to a single robot sampling approach (Manjanna et al., 2018). We do not formulate the multi-robot scenario as a distributed or collective learning problem. Instead, we train each robot as an independent agent¹ and observe their sampling performance when they

This is one of the several papers published in *Autonomous Robots* comprising the Special Issue on Robot Swarms in the Real World: from Design to Deployment.

✉ Sandeep Manjanna
msandeep@seas.upenn.edu; msandeep.sjce@gmail.com

M. Ani Hsieh
m.hsieh@seas.upenn.edu

Gregory Dudek
dudek@cim.mcgill.ca

¹ GRASP Laboratory, University of Pennsylvania, Philadelphia, PA, USA

² Center for Intelligent Machines, McGill University, Montreal, QC, Canada

¹ The robots are assumed to be homogeneous in terms of their capabilities. They may all have the same set of learnt parameters. However, this is not compulsory as long as they are all trained on similar distributions.




Fig. 1 Aerial view of two robotic boats surveying the reef at a marine reserve off the coast of Barbados island. Example paths of the boats are illustrated with solid lines of colors red and white. The goal is to plan paths that are rewarding in terms of information gain and have minimal overlap. We achieve higher rewards by determining the best policy and minimal overlap through communication between the robots




Fig. 2 Overview of our approach. We train each individual robot to perform efficient sampling and deploy them together to analyze the coverage behavior when the robots are given an ability to communicate

are put together in a sampling task and are given the ability to intermittently communicate as shown in Fig. 2. Doing so makes the problem more tractable; and the multi-robot system becomes easily scalable as the same trained parameters can be used on any new robot added to the system. The team becomes more resilient to failure as all robots are individually trained to perform the task efficiently, and the team can be put together in any robot combination.

An overview of the training and testing phases of an individual robot is illustrated with a block diagram in Fig. 3. In the training phase, the parameters for the policy are learnt on a set of generic distributions that represent the field that needs to be surveyed. Then in the testing phase, when the robot is being deployed in the field, the prior of the current field is used to generate a policy for the robot given its current location or state. In our example application of sampling visual data of the coral reef using a robotic boat, we train the algorithm with a set of satellite images of different coral regions. While deploying the robot with these learnt parameters from the training, we use an aerial image of the survey area to provide a prior scoremap for the algorithm to generate an action plan for the robotic boat.

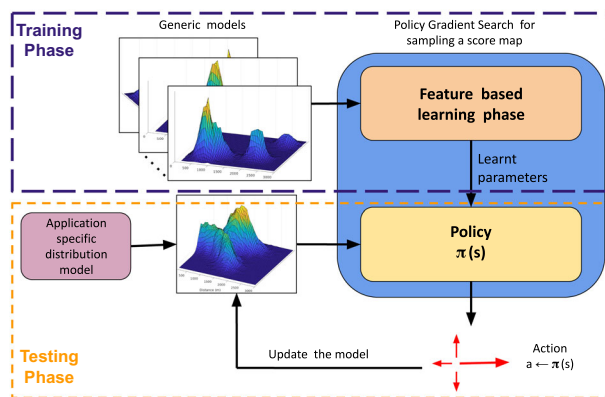


Fig. 3 Overview of the training and testing phases

Once the robots are trained, they are all put together in a sampling task and are allowed to communicate only their visited locations with other neighboring robots that are within the communication range. Communication between agents plays an important role in a decentralized multi-robot system. In this paper, we consider the effects of communication between multiple robots on the overall sampling performance of the team. We compare the overall sampling performance of the team as the communication range of the robots changes. Our proposed multi-robot sampling algorithm performs efficiently even when robots in the team are operating with a limited communication range, thus demonstrating the scalability of our approach for sampling large-scale environments.

Key contributions of this paper include:

- A scalable multi-robot system with distributed decision making to achieve nonuniform sampling of a quasistatic field.
- Analysis of the effect of robot communication range on the overall sampling performance of the team.
- Investigation of the task-sharing behavior between robots.

We present an empirical evaluation of our sampling technique with statistically significant results. We demonstrate our methodology with sea-floor bathymetry data collected during our field trial as illustrated in Fig. 4. We also evaluate our sampling approach on a set of diverse environments presented in Fig. 5, some synthetically generated and others commonly occurring in nature. We compare our sampling algorithm to baseline strategies through simulated experiments that utilize models derived from both synthetic and real robot deployment data.

2 Related work

There is a large literature on robotic data sampling and coverage. Especially in the marine domain, much of this is based on

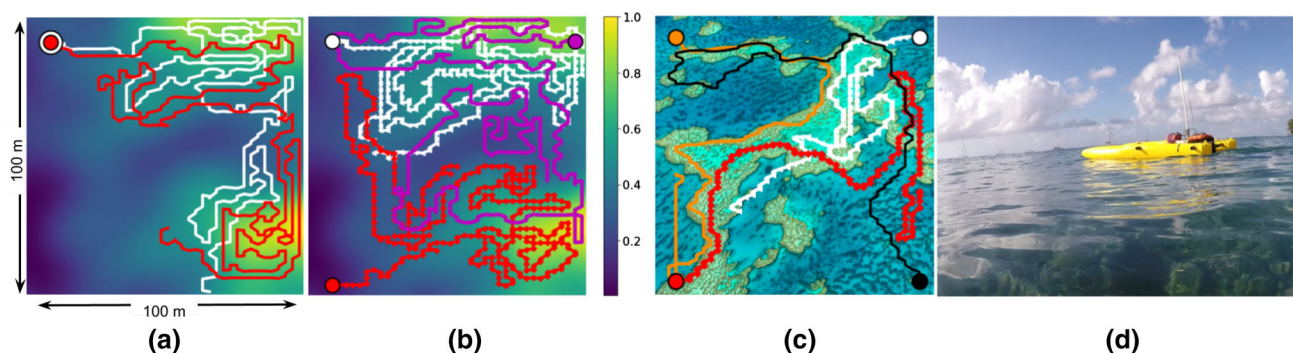


Fig. 4 **a** Two-robot and **b** Three-robot sampling paths planned over bathymetry data collected during our field trials in Barbados. **c** Four-robot sampling paths planned over a satellite image of a coral reef.

Colored circles represent the start points of the robots. **d** An autonomous surface vehicle (ASV) deployed to sample the bathymetry and visual data over the reefs in Barbados




Fig. 5 Distributions used for analyzing the multi-robot policy gradient spatial sampling algorithm. The first two distributions (**a**, **b**) are generated synthetically using a mixture of Gaussians. Distribution in **c** is a scoremap generated using real reef data (multi-spectral processing of an

aerial image over a reef). **d** is the diffusive spreading pattern simulating the impacts of different nutrients on algae (Dai & Zhao, 2014). **e** is a trimmed section of the distribution in **d**

mechanisms that use waypoints of geometric priors with limited dependence on the distribution of incoming observations (Shkurti et al., 2012; Bender et al., 2013; Stephanie Kemna et al., 2018; Manjanna & Dudek, 2017). Many techniques use data regarding boundaries of the environment (Choset & Burdick, 1996) and some methods try to minimize odometry error or energy utilization (Sipahioglu et al., 2010; Xu et al., 2011). Our work places emphasis on the density of valuable measurements that can be collected in a limited amount of time.

The design of pragmatic, efficient coordinated multi-robot exploration algorithms to carry out the mission reliably and quickly is a topic of active research. The use of multiple robots instead of a single robot is often suggested to have several advantages and leads to a variety of design trade-offs. Classical algorithmic results related to computational complexity perform poorly in stochastic and time-varying real-world scenarios (Dudek & Roy, 1997; Yan et al., 2013; Higuera & Dudek, 2013). Another advantage of robot teams is due to the merging of overlapping information, which can

help compensate for sensor uncertainty (Burgard et al., 2005; Almadhoun et al., 2019).

Robotic sampling requires an efficient strategy to decide where to sample so that there is the highest knowledge gain. Decision making can be regarded as a cognitive process resulting in the selection of one or more actions (i.e., a policy) among several alternatives. Two popular approaches for decision making in multi-robot scenarios are centralized and decentralized. A decentralized approach for decision making is scalable, robust, and efficient (Dudek et al., 1996; Vaughan et al., 2002). Generally, the problem of distributed multi-robot exploration can be stated as n identical robots set out to explore an unknown area, each robot is equipped with sensing, localization, mapping, and limited-range communication capability (Sheng et al., 2006). Julian et al. propose an information theoretic approach to distributively control multiple robots equipped with sensors to infer the state of an environment (Julian et al., 2012). In their approach, the robots iteratively estimate the environment state using a sequential Bayesian filter, while continuously moving along

the gradient of mutual information to maximize the informativeness of the observations provided by their sensors. However, they assume a connected communication graph throughout the exploration task. Instead, in our approach, we consider a fixed distance communication disk model around the robot (Spanos & Murray, 2005; Ji & Egerstedt, 2007; Zavlanos & Pappas, 2008). This allows for intermittent communication between the robots, which is more realistic in large-scale applications.

Even though multi-robot systems have numerous advantages, they do have additional costs. One such overhead is coordination and communication between the robots. The overall system performance can be directly affected by the quality of coordination and control. Communication, as a means of coordination, enables robots to share position information, state of the environment, sensor measurements, and enable individual robots to learn about the intentions, goals, and actions of other robots. Yan et al., in their survey, classified the communication structure based on the information transfer modes, such as explicit and implicit communication (Yan et al., 2013). Explicit communication allows direct exchange of information between the robots (Freda et al., 2019; Lowe et al., 2017; Salam & Hsieh, 2019), whereas in implicit communication, robots get information about their fellow robots through the environment (Godoy et al., 2016). Explicit communication model enables the robots to work in a team towards completing the task efficiently. In this paper, we explore explicit and asynchronous communication between robots with a constraint on the communication range.

Compared to complete coverage algorithms, adaptive sampling approaches trade off completeness for efficiency. In many cases, even if the process being studied is rapidly varying, subsampling can be effective when the sample points are correctly selected (Venkataramani & Bresler, 2000). When the environmental phenomena being sampled are smoothly varying without any local maxima peaks, non-adaptive strategies are known to perform well (Singh et al., 2006). However, if the environment contains peaks with high local-variance, adaptive sampling can exploit the clustering phenomena to map the environmental field more accurately than non-adaptive sampling (Manjanna & Dudek, 2017). In our multi-robot adaptive sampling approach, we explore the possibility of multiple independently trained robots performing the sampling task efficiently with minimal communication between the robots. Our formulation does not encompass the objective of minimizing the communication, but it is a by-product of our sampling approach. We compare the overall sampling performance of the team by varying the communication range of the robots, thus evaluating our approach for scalability in large-scale sampling problems.

3 Problem formulation

The sampling region is a continuous two-dimensional area of interest $\mathcal{E} \subset \mathbb{R}^2$ with user-defined boundaries. The spatial sampling region is discretized into uniform grid cells, such that the k th robot's position $\bar{\mathbf{x}}_k$ can be represented by a pair of integers $\bar{\mathbf{x}}_k \in \mathbb{Z}^2$. Each grid cell (i, j) is assigned a prior score value $q(i, j)$ of the data that can be sampled in that cell. We assume that a low-quality initial estimate of the phenomenon being sampled is known either through pilot surveys or satellite data and use this prior to initialize the grid cell scores.

The objective is to maximize the total accumulated score by all K robots, $J_{total} = \sum_{k=1}^K J_k$. This in turn can be achieved by maximizing the accumulated score by each robot J_k . For brevity, we will use J to indicate each robot's accumulated score. Thus, the goal of an individual robot is to maximize J over a trajectory τ within a fixed amount of time T . We will specify each robot's behavior using a parameterized policy. This is a conditional probability distribution $\pi_{\bar{\theta}}(\bar{\mathbf{s}}, \bar{\mathbf{a}}) = p(\bar{\mathbf{a}}|\bar{\mathbf{s}}; \bar{\theta})$ that maps the current state $\bar{\mathbf{s}}$ of the survey to a distribution over possible actions $\bar{\mathbf{a}}$. Our aim will be to automatically find good parameters $\bar{\theta}$, after which the policy can be deployed without additional training on sampling new environmental fields.

Finding a sequence of actions that maximizes a long-term objective could be done using dynamic programming or other iterative techniques. However, in our formulation, the system state is described using a map containing the per-cell score or value of the physical phenomenon being measured and this scoremap changes as areas are visited by the robot. Hence, the state is composed of the agent's current position and the scoremap containing the per-location value of the phenomenon being measured. As a result of formulating complex states, the state space grows exponentially with the size of the sampling region. Dynamic programming is impracticable with such a large state space—especially if the time to solve each particular problem is limited. Instead, we turn to methods that directly optimize the policy parameters $\bar{\theta}$ based on (simulated) experiences. To apply these methods, we will first formalize the sampling problem as a Markov Decision Process (MDP).

3.1 Formalizing sampling as an MDP

An MDP is a formal definition of a decision problem that is represented as a tuple $(S, A, T(\bar{\mathbf{s}}_{t+1}|\bar{\mathbf{s}}_t, \bar{\mathbf{a}}_t), r(\bar{\mathbf{s}}_t, \bar{\mathbf{a}}_t), \gamma)$, where S and A are the state and action space, T models transition dynamics based on the current state and action, and r defines the reward for the current state-action pair. γ is a discount factor that reduces the desirability of obtaining a reward t time-steps from now rather than in the present by γ^t . The objective is then to optimize the expected discounted

cumulative reward $J = \mathbb{E}_\tau[\sum_{t=0}^H \gamma^t r_t(s_t, a_t)]$, where H is the optimization horizon.

Our current sampling formulation considers the state \vec{s} to include both position of the k^{th} robot \vec{x}_k as well as the map q containing per-location utility value of the phenomenon being measured, $\vec{s} = (\vec{x}_k, q)$. To begin with, we consider a 4-connected action space with the options for the robot to move to and scan the cell North, East, South or West of its current location. The action space can be easily expanded or constrained according to the motion constraints of the platform used. We consider a probabilistic choice of actions, that is, an action is chosen according to the policy distribution $\pi_\theta(\vec{s}, \vec{a})$. Once the data in the current robot location is measured, the utility value $q(i, j)$ of the current grid-cell (i, j) is reduced to 0. The discounted reward function is defined as $\gamma^t q(\vec{x}_k)$, with the discount factor $0 \leq \gamma \leq 1$ encouraging the robot to sample cells with high scores in early time steps t . Discount factor γ with a value closer to 0 encourages myopic trajectories, whereas a γ closer to 1 pushes towards achieving trajectories that are farsighted (meaning paths that achieve higher accumulated rewards)

4 Policy gradient for sampling

The complex and large state space required to formulate the sampling problem makes it unrealistic to apply dynamic programming or other iterative techniques for robotic sampling problems. Even though we have used iterative methods to solve similar problems, they do not scale well with the increase in the size of the sampling region (Manjanna & Dudek, 2017). Hence, we make use of policy gradient search that directly optimizes the policy parameters θ based on (simulated) experiences.

Policy gradient methods use gradient ascent for maximizing the expected return J_θ . The gradient of the expected return ($\nabla_\theta J_\theta$) guides the direction of the parameter (θ) update. The policy gradient update is given by,

$$\theta_{k+1} = \theta_k + \eta \nabla_\theta J_\theta, \tag{1}$$

where η is the learning rate. The policy gradient is given by,

$$\nabla_\theta J_\theta = \int_\tau \nabla_\theta p_\theta(\tau) R(\tau) d\tau, \tag{2}$$

where $R(\tau)$ is the reward obtained by following the trajectory τ . One of the effective methods to estimate the gradient $\nabla_\theta J_\theta$ is to make use of the *likelihood-ratio* trick that is given by the identity $\nabla p_\theta(y) = p_\theta(y) \nabla \log p_\theta(y)$. Applying this likelihood-ratio trick to the policy gradient in Eq. 2, decomposing $\nabla_\theta \log p_\theta(\tau)$ into single time steps and applying the

logarithm yields,

$$\nabla_\theta \log p_\theta(\tau) = \sum_{t=0}^{H-1} \nabla_\theta \log \pi_\theta(a_t | s_t, t) \tag{3}$$

One of the first policy gradient algorithms, REINFORCE Algorithm (Williams, 1992), applies Eq. (3) and uses a variance-reducing baseline b to define the policy gradient as,

$$\nabla_\theta J_\theta = \mathbb{E}_{p_\theta(\tau)} \left[\sum_{t=0}^{H-1} \nabla_\theta \log \pi_\theta(a_t | s_t, t) (R(\tau) - b) \right] \tag{4}$$

This expression, however, depends on the correlation between actions and previous rewards. These terms are zero in expectation, but can often induce additional variance. Ignoring these terms yields lower-variance updates. G(PO)MDP algorithm (Greensmith et al., 2004) and Policy Gradient Theorem (PGT) algorithm (Sutton et al., 2000) propose to reduce the variance of policy gradient estimates by using the observation that rewards from the past do not depend on actions in the future. Accordingly, the policy gradient is given by,

$$\nabla_\theta J_\theta = \frac{1}{m} \sum_{i=1}^m \sum_{t=0}^{H-1} \nabla_\theta \log \pi_\theta(a_t^{(i)} | s_t^{(i)}) \left(\sum_{j=t}^{H-1} r(s_j^{(i)}, a_j^{(i)}) - b(s_t^{(i)}) \right). \tag{5}$$

The expectation over $p_\theta(\tau)$ is approximated by summing over all the sampled trajectories ($\tau^{(i)} = (s_0^{(i)}, a_0^{(i)}, s_1^{(i)}, a_1^{(i)}, \dots)$). The gradient is based on m sampled trajectories from the system, with $\vec{s}_j^{(i)}$ the state at the j th time-step of the i th sampled roll-outs. Furthermore, b is a variance-reducing baseline. In our experiments, we use the observed average reward as baseline ($b = \mathbb{E}[R(\tau)] \approx \frac{1}{m} \sum_{i=1}^m R(\tau^{(i)})$).

4.1 Policy design

An effective approach to define stochastic policies over a set of deterministic actions is the use of Gibbs distribution of a linear combination of features as a policy (also referred to as Boltzmann exploration of the softmax policy). We consider a commonly used linear Gibbs softmax policy parametrization (Sutton et al., 2000) given by,

$$\pi(s, \mathbf{a}) = \frac{e^{\theta^T \phi_{sa}}}{\sum_b e^{\theta^T \phi_{sb}}}, \quad \forall s \in S; a, b \in A, \tag{6}$$

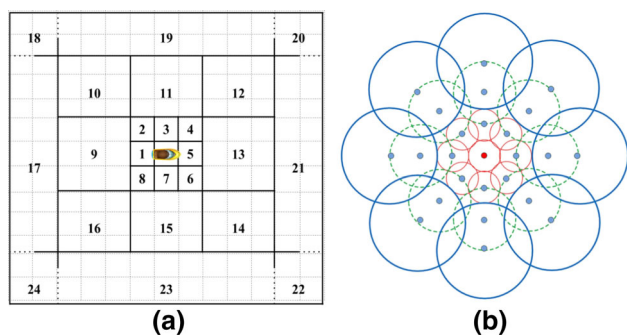


Fig. 6 **a** Multiresolution aggregation of feature space. **b** DAISY: local image descriptors that inspired our feature aggregation (Tola et al., 2009)

where ϕ_{sa} is a k -dimensional feature vector characterizing a state-action pair (s, a) and θ is an k -dimensional parameter vector. This is a commonly used policy in reinforcement learning approaches.

The final feature vector ϕ_{sa} is formed by concatenating a vector $\phi'_s \delta_{aa'}$ for every action $a' \in [North, East, South, West]$, where $\phi'_s \in \mathbb{R}^k$ is a feature representation of the state space, and $\delta_{aa'}$ is the Kronecker delta. Thus, the final feature vector has $4 \times k$ entries, 75% of which corresponding to non-chosen actions will be 0 at any one time step.

4.2 Feature design

In our previous work (Manjanna et al., 2018), we proposed and evaluated multiple feature aggregation designs. Our empirical analysis showed that having a multiresolution aggregation of the feature space resulted in achieving better discounted rewards. In multiresolution feature aggregation, feature cells grow in size along with the distance from robot location as depicted in Fig. 6a. Thus, areas close to the robot are represented with high resolution and areas further from the robot are represented in lower resolution. We draw our inspiration for this feature design from a popular local image descriptor called DAISY descriptor (Tola et al., 2009) presented in Fig. 6b. The intuition behind this feature design is that the location of nearby reward values is important to know exactly, while the location of faraway rewards can be represented more coarsely. The multiresolution feature design is also suitable for bigger worlds as it scales logarithmically with the size of the world.

5 Sampling with multi-robot teams

We propose a decentralized sampling approach where each robot in a team performs an informed survey using a policy-gradient-based sampling strategy as outlined above. As mentioned earlier, we do not formulate the multi-robot sce-

nario as a learning problem. Instead, we train each robot as an independent agent.

One way to achieve collaborative sampling between multiple independent robots is by assigning a reward function for the robot r_k that is proportional to the distance of grid location l_i from other $n - 1$ coworking robots and inversely proportional to the distance of grid location l_i from its current location l_{r_k} . Thus encouraging robot k to visit locations that are close to the l_{r_k} and farther from other co-working agents. This results in dividing the task spatially so that each robot has its own subregion to sample.

$$R(l_{r_k}, l_i) \propto \frac{\sum_{j \in (1, 2, \dots, n), j \neq k} \text{distance}(l_{r_j}, l_i)}{\text{distance}(l_k, l_i)}. \quad (7)$$

However, performing such weighted reward adjustments can compromise the resilience and robustness of the multi-robot team. Overall, it will degrade the sampling performance as the robots can and do get stuck within their own self-defined region of space. Instead, in our proposed approach, the robots communicate only their visited locations with the neighboring robots within their communication range. This information is used by each robot to transition to the state where the rewards from all the visited locations are removed. In this work, we explore how different communication ranges impact the overall sampling performance of the robot team. For our empirical analysis, we run experiments on the distributions presented in Fig. 5.

5.1 Distance constrained communication

We explore the ability of a multi-robot team to sample efficiently with a distance-limited communication between the robots. Every robot can communicate only its current state with the neighboring robots that are within the communication range. This state information is then used by neighboring robots to update their sampling plan accordingly.

This setup with distance constrained communication imitates the real-world scenario in any outdoor deployment environment. Especially on the surface of water, the communication signals attenuate faster (Coelho et al., 2018) making it infeasible to exchange messages between robots that are far apart. We define the distance constraints as a percentage of the maximum possible distance (D_{max}) in the bounded survey region. Thus, 0% represents no communication and 100% represents complete communication throughout the survey region.

We note that robots can choose the same overlapping path when there is no communication between them under this strategy. Sometimes if the robots are synchronized in sampling, they will end up following the same path even after communicating on every step as they have knowledge only

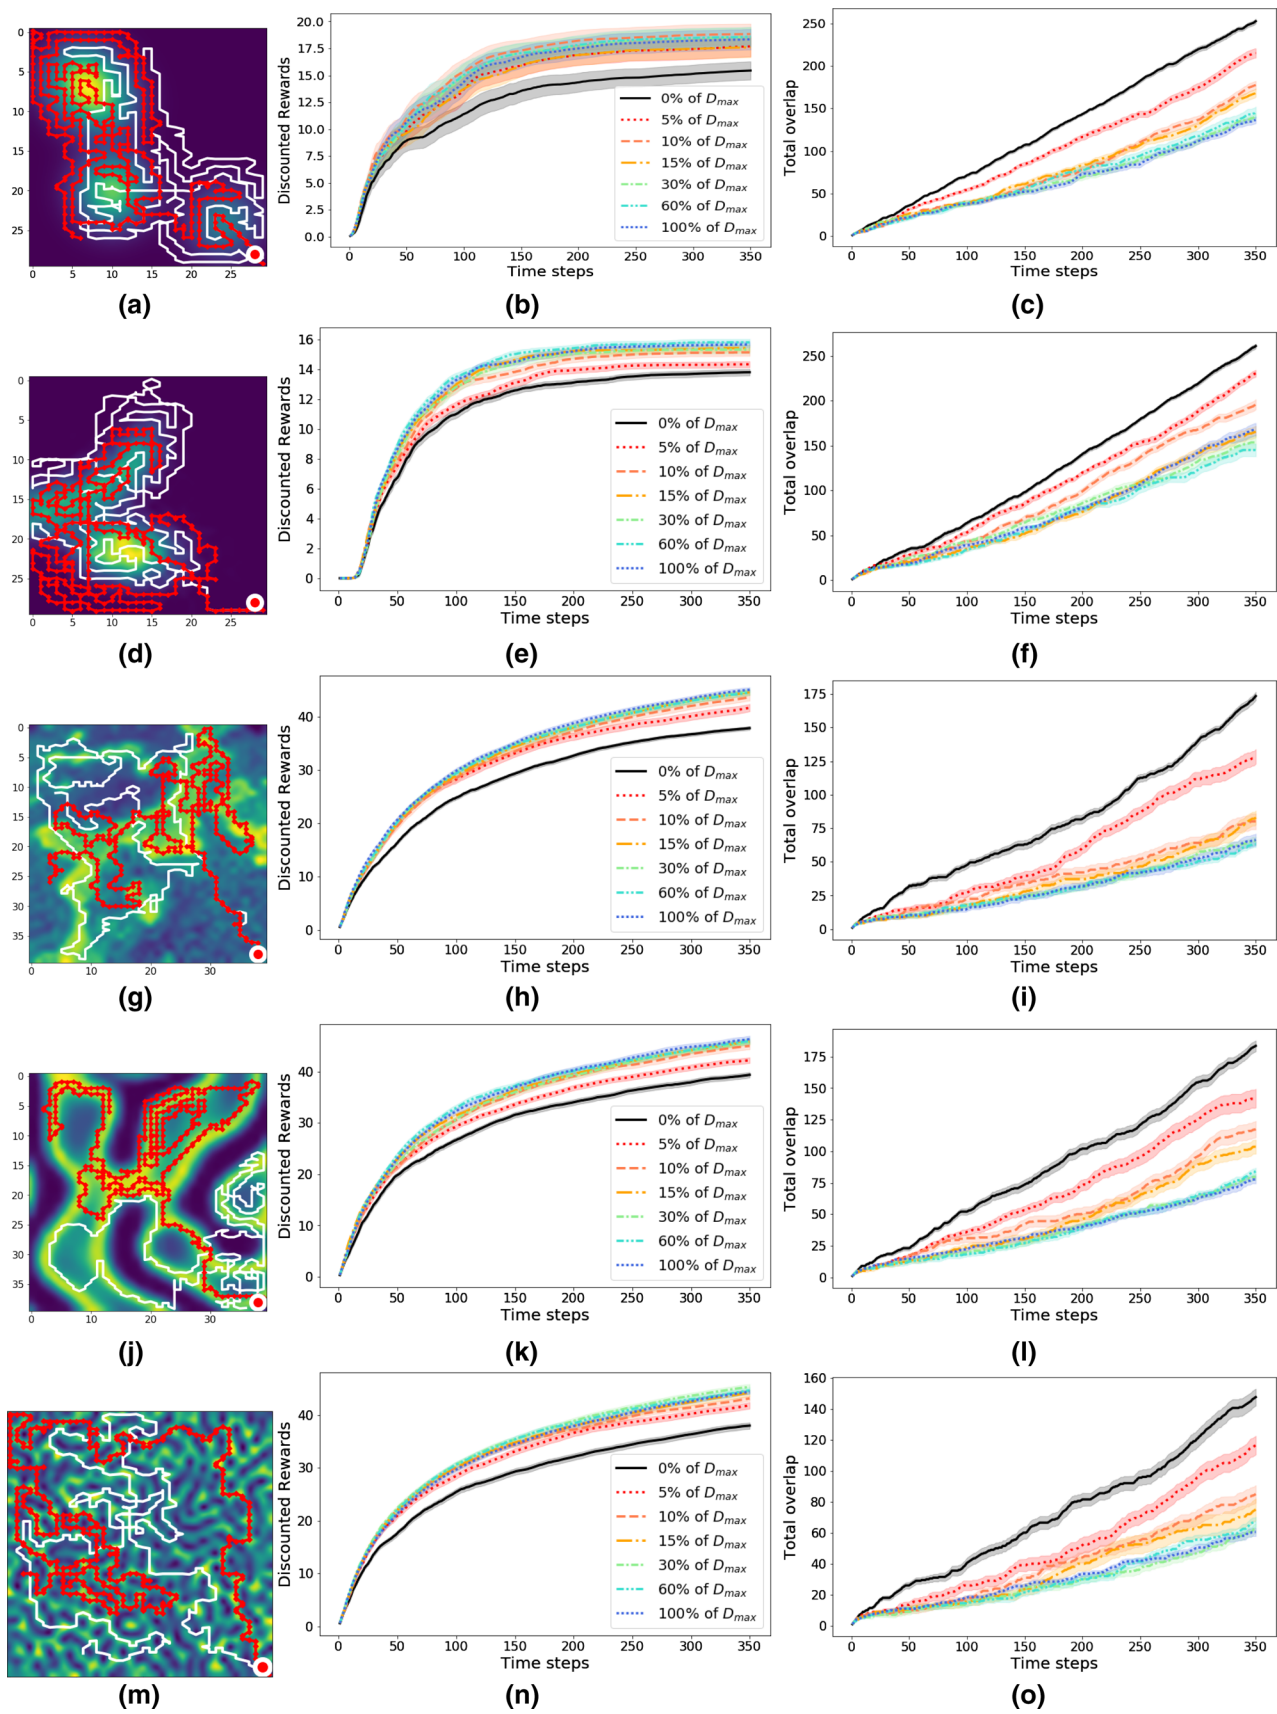


Fig. 7 Sampling performance and illustrative paths of two robots operating with distance constrained communication. The shaded regions around the median line illustrate the standard error over 20 trials for each communication range

about the current step. Moreover, having such frequent communication is expensive and infeasible in the real world. Our previous observations show that a communication after the collection of every 10–30 samples would fetch good rewards.

6 Experimental results

6.1 Setup

We train all our robots on a set of distributions generated by a mixture of Gaussians. All robots have identical learnt parameters, thus are independently capable of collecting good samples in the data collection task. A diverse set of five distributions is used to simulate the physical phenomenon of interest. Two of these distributions are synthetically generated using a mixture of Gaussians (Fig. 5a, b). The next distribution in Fig. 5c is a scoremap generated by multispectral image processing on an aerial image of the coral reef. Last two distributions in Fig. 5d, e are generated using the model for a diffusive spreading pattern simulating the impacts of different nutrients on algae (Dai & Zhao, 2014). Over the multiple trials, we combine different starting states of the robots (such as all robots starting at a fixed location, or randomly chosen location, or every robot starting at opposite corners). We experimented with 10 communication ranges. In this section, we only present a small set of our experimental results to achieve better readability and clarity in the plots. We apply Savitzky–Golay filter (Savitzky & Golay, 1964) to smoothen the planned paths. This kind of path smoothing is used only if the paths are planned offline. Many applications in the domain of environmental monitoring with autonomous vehicles have a need to sample in information rich locations at the earliest (Low et al., 2008). This is true when robotic vehicles have a limited endurance to carry out the survey and the environmental processes are quasistatic. Hence, we measure the total accumulated discounted rewards as one of the performance metrics. In our experiments we use hyperbolic discounting for training as it encourages a non-myopic reward collection by the agent. Two or more robots surveying a region to collect samples are expected to cover different regions so that the resources are utilized efficiently. Hence, we measure the overlap between the robot paths as the second metric for sampling performance. To achieve an efficient spatial sampling, we want the planned paths to achieve as little overlap as possible and as high discounted rewards as possible.

6.2 Results and discussion

Figure 7 demonstrates the results from our experiments involving two robots starting their sampling path from the same starting point. The first column in Fig. 7 illustrates the

example paths generated by two robots with full communication. The second and the third columns in Fig. 7 present the discounted rewards and the path overlap, respectively. The shaded area specifies the standard error of the mean over 20 trials for each of the 7 communication ranges. These results illustrate that, irrespective of the underlying distribution, a communication range between 10% and 30% of the D_{\max} can achieve good sampling with minimal message exchange. To confirm this hypothesis, we further analyze the results by averaging over all 5 spatial distributions.

Figure 8a presents the averaged results over all spatial distributions plotted against the communication range on the x-axis. This plot indicates a significant improvement in the reward collection and a significant decrease in the path overlap once the communication range is over 15% of the maximum possible distance D_{\max} . Another observation from Fig. 8a is that the robots are achieving good sampling performance with minimal message exchange between them. This resulting behavior of increased sampling performance with minimal communication is a desired feature for a decentralized multi-robot sampling approach.

We evaluate the scalability of our approach by increasing the number of robots involved in sampling (Fig. 8c). The plots presented in Fig. 8b illustrate the change in total rewards collected by the robotic team as the number of deployed robots increases (x-axis). It can be observed that the curves for the total rewards collected plateaued with the increase in the number of robots. This is because of the bounded sampling region where all robots are fenced according to our formulation of the problem.

We also evaluated the robot sampling technique on real bathymetry data collected using two ASVs (Fig. 4d) during our field trip in Barbados in 2019. Fig. 4a, b present the paths of multiple robots sampling from these depth maps, where larger rewards are associated with shallower regions, thus inducing the robots to cover relatively shallow regions first. Figure 4c demonstrates the path planning performed on an aerial image of the coral reef. We demonstrate the resulting paths when a varying number of robots start at different locations, as illustrated in Fig. 4. These results establish the applicability of our distributed sampling technique for real-world data sampling scenarios.

Experiments in Fig. 4a present the sampling of bathymetry data of the seafloor using two autonomous surface vehicles. Water in its pure form is an insulator, but as found in its natural state, it contains dissolved salts and other matter which makes it a partial conductor. The higher its conductivity, the greater the attenuation of radio signals which pass through it VK5BR (1987). In a recent study (Dala & Arslan, 2021), a Long Range (LoRa) communication antenna was found to achieve a communication range between 80 and 160 m depending on whether that antenna was buffering or not. Using our proposed algorithm, a region of 1×1 km can

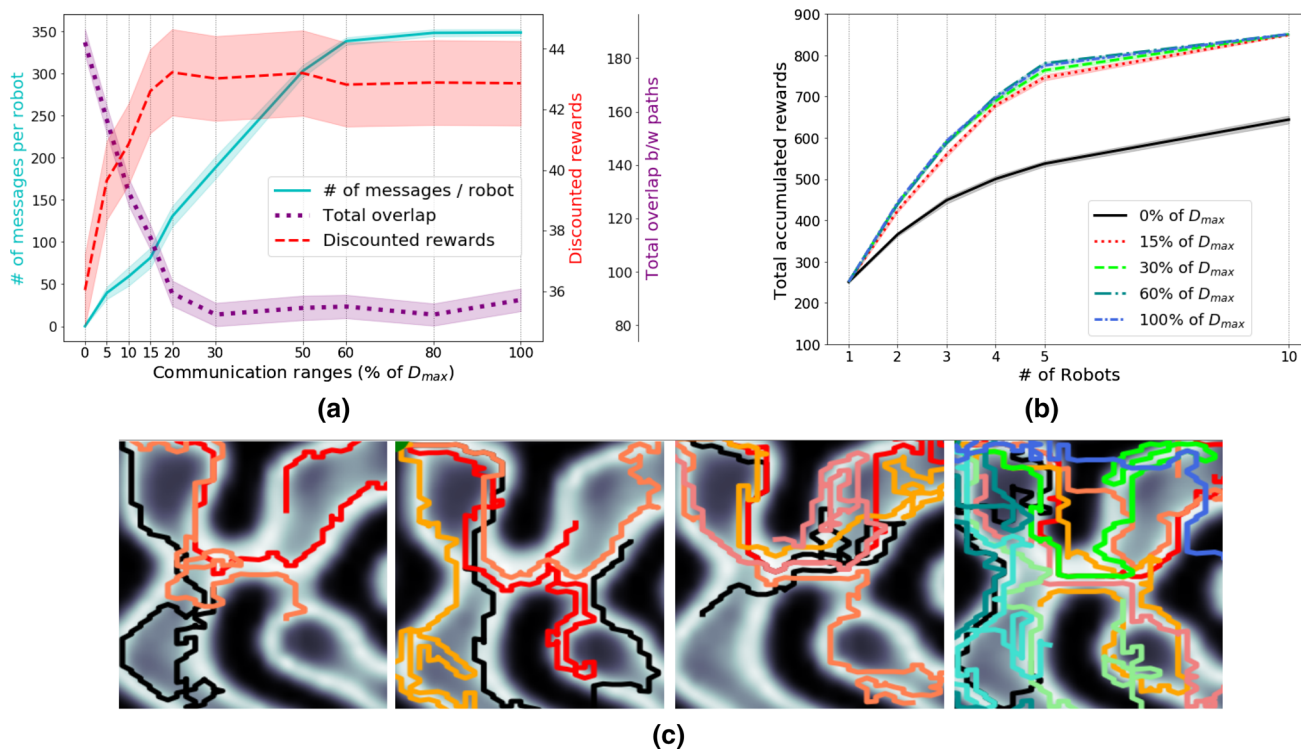


Fig. 8 **a** Accumulated rewards as the number of robots is increased. The shaded area represents the standard error over 20 trials. **b** Sampling performance metrics averaged over all the 5 distributions. These results illustrate higher sampling performance achieved even with mini-

mal communication. The shaded area represents the standard error over 100 trials. **c** Paths for varying team size (2, 4, 5, and 10) on an example static spatial field

be efficiently sampled by two robots with a communication range of 140m, which is a realistic range on the surface of water to achieve reliable communication connection.

6.3 Performance comparison with baseline techniques

We compare our proposed algorithm to two baseline sampling strategies through simulated experiments that utilize models derived from both synthetic and real robot deployment data. The first comparison algorithm is a multi-robot coverage path planning technique called DARP algorithm (Divide Area based on the robot’s initial Positions) Kapoutsis et al. (2017). DARP algorithm divides the region into a number of equal areas each corresponding to a specific robot, such that a complete coverage is guaranteed, the coverage path is of minimum cost and non-backtracking. The second comparison algorithm is a multi-robot adaptation of the maxima search algorithm (Meghjani et al., 2016). Each robot sequentially selects the current maxima in the given spatial field and plans a path to that location using A-star path planner. Once the value at a given location is collected, it is reduced to 0 as in our proposed approach. The first three columns

of Fig. 9 present the illustrations of fixed horizon paths executed by the two robots using our proposed approach, the DARP algorithm, and the multi-robot maxima search algorithm, respectively. The comparison of discounted rewards achieved by all three techniques is presented in the fourth column of Fig. 9.

We evaluate our proposed approach by comparing the performance in terms of total accumulated rewards and the total discounted rewards. We found that our proposed policy gradient based sampling technique performs better than the other two baseline approaches by collecting high rewarding samples in the early phase of their surveys. As illustrated in Fig. 9d, h, our approach performs significantly better for sampling fields with uneven hotspot distribution. If the process of interest is uniformly distributed as in Fig. 9i, our approach performs comparably with the complete coverage approach (Fig. 9l). Many environmental processes have a nonuniform spatial distribution as presented in the depthmaps of the coral reefs presented in Fig. 9m and our proposed approach was able to collect high rewarding samples at the early phase of the survey as depicted by the results in Fig. 9p.

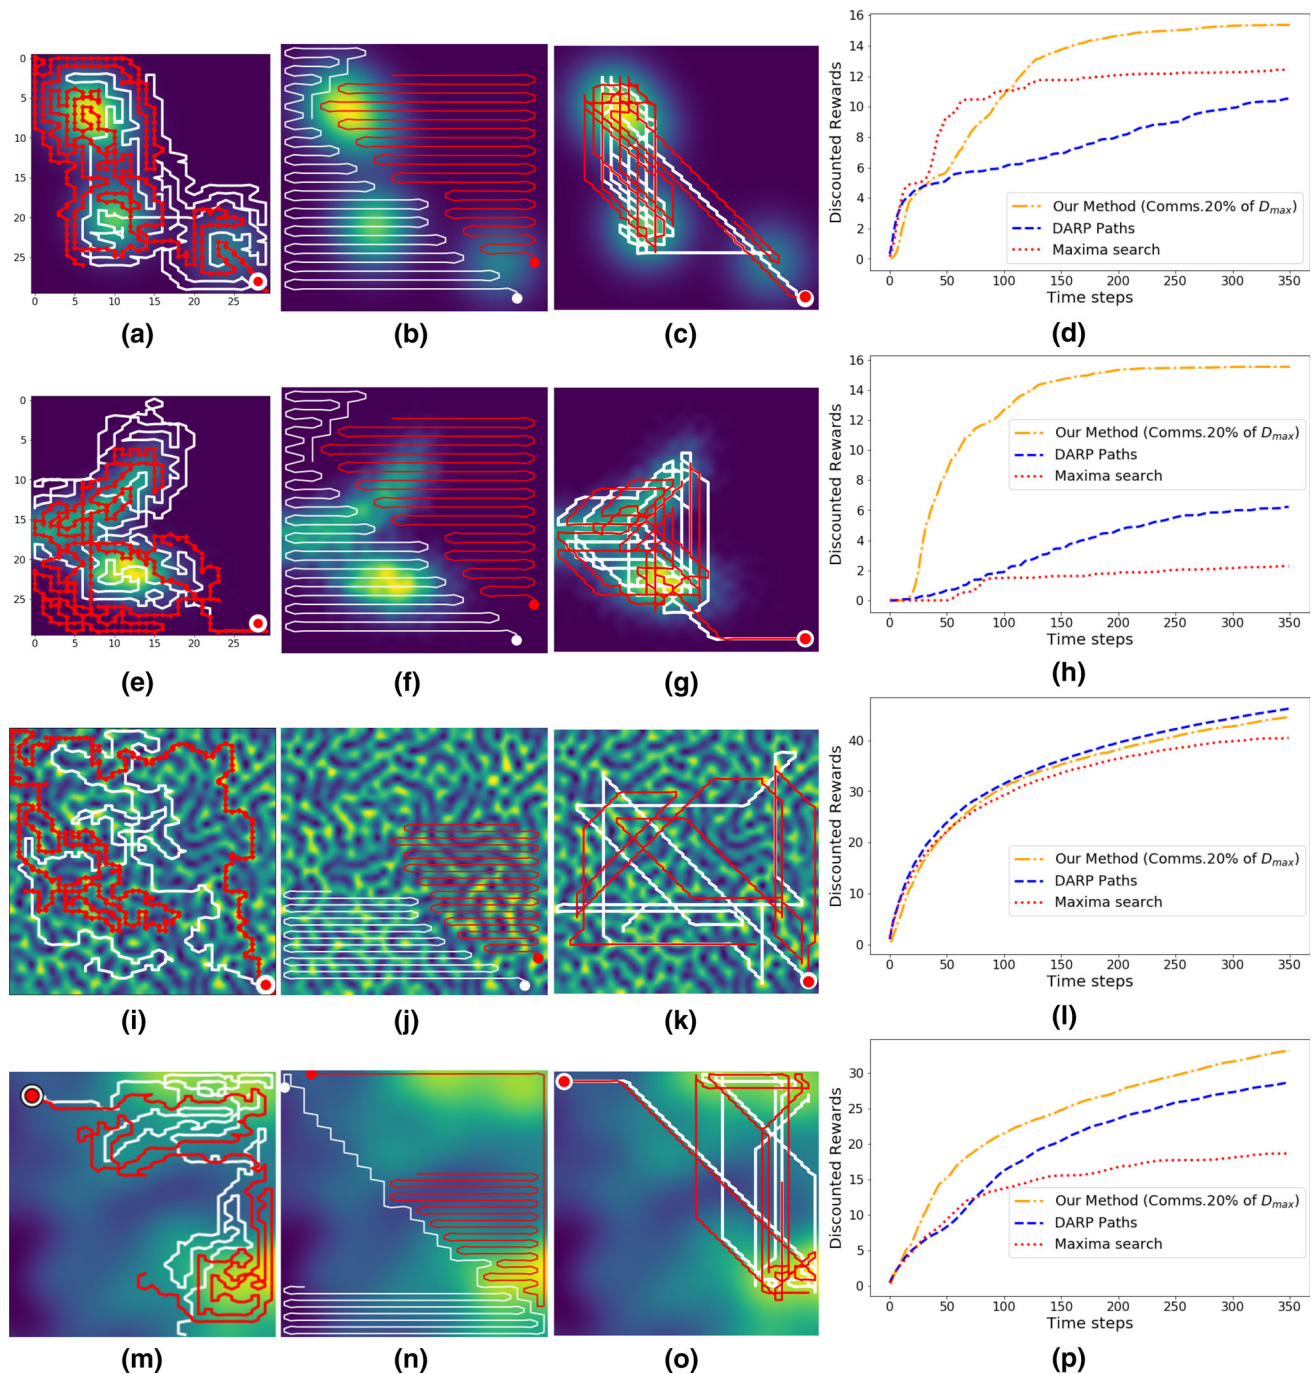


Fig. 9 Comparing the performance of our sampling approach with two baseline techniques. Illustrative paths of two robots following our algorithm and the baseline algorithms

6.4 Additional properties

In this section, we present additional properties and potential extensions to the proposed planning algorithm. The results in Fig. 10b illustrate an example of the robustness of our multi-robot sampling system. As shown in Fig. 10a, when there are no failures, both the robots share and cover the hotspot

regions. The robots are all individually trained to perform efficient sampling of a spatial field. Hence, even when one of the robots fails during a survey, the other robot samples from the hotspot regions and achieves a good overall performance for the team as depicted in Fig. 10b. These preliminary results display a fault-tolerance nature of our multi-robot sampling




Fig. 10 **a** Sampling paths of the robots when there are no failures. **b** Sampling paths of the robots when one of the robots fail after first 75 time steps. **c** Paths planned for two robots sampling with a constrained action space. Colored stars represent the start points of the robots

strategy and we would like to further investigate this behavior in the near future.

In Sect. 3.1, we described that the problem formulation and our solution can be easily extended to any kind of action space. The preliminary results validating these claims are illustrated in Fig. 10c. Here, we present an example of two robots operating in constrained action space with a suite of only three “forward moving” actions available in any state: turn 45° left, go straight, and turn 45° right. All these actions are in relation to the current heading of the robot and hence the heading should be included into our state space to plan for this action space. We are working on extending our planner to plan over action spaces designed using motion primitives that are specific to a robotic platform.

7 Conclusions and future directions

In this paper, we presented a distributed path planning approach to generate an efficient sampling path for a team of robots. Our objective is to have paths that are highly rewarding in terms of information gain and have minimal overlap. We achieve higher rewards by using data-driven methods to determine the best policy and attain minimal overlap through communication between the robots. Instead of formulating the multi-robot scenario as a learning problem, each robot is trained independently and put together to perform a sampling task. Our analysis of the coordinated sampling behavior shows that the robot team was able to efficiently collect high rewards with a minimal exchange of messages between the robots.

Our results demonstrate that, irrespective of the underlying distribution that is being sampled, our approach is able to collect high rewards with a limited communication range of 10–20% of the D_{\max} . These outcomes are significant when sampling in large-scale environments where a complete communication between the robots is unrealistic.

We also empirically analyzed the scalability and adaptability of our approach with changing the number of robots and changing action space.

In the ongoing work, we will be examining a larger scale deployment of robots using the methods examined here, and the impact of physical disturbances on the system. In the near future, we plan to theoretically analyze the coordination behavior of the sampling robots. One interesting idea would be to learn the number of robots needed to efficiently sample a given spatial field using prior knowledge about the field itself.

References

- Almadhoun, R., Taha, T., Seneviratne, L., & Zweiri, Y. (2019). A survey on multi-robot coverage path planning for model reconstruction and mapping. *SN Applied Sciences*, 1(8), 1–24.
- Bender, A., Williams, S. B., & Pizarro, O. (2013). Autonomous exploration of large-scale benthic environments. In *IEEE international conference on robotics and automation* (pp. 390–396).
- Burgard, W., Moors, M., Stachniss, C., & Schneider, F. E. (2005). Coordinated multi-robot exploration. *IEEE Transactions on Robotics*, 21(3), 376–386.
- Choset, H., & Burdick, J. (1996). Sensor based motion planning: The hierarchical generalized Voronoi graph. In *Algorithms for robot motion and manipulation* (pp. 47–61).
- Coelho, A., Lopes, M., Ferreira, B., Campos, R., & Ricardo, M. (2018). Experimental evaluation of shore to unmanned surface vehicle WI-FI communications. In *2018 wireless days (WD)* (pp. 86–91). IEEE.
- Dai, C., & Zhao, M. (2014). Nonlinear analysis in a nutrient-algae-zooplankton system with sinking of algae. In *Abstract and applied analysis* (Vol. 2014). Hindawi.
- Dala, A., & Arslan, T. (2021). Design, implementation, and measurement procedure of underwater and water surface antenna for Lora communication. *Sensors*, 21(4), 1337.
- Dudek, G., Jenkin, M. R., Milius, E., & Wilkes, D. (1996). A taxonomy for multi-agent robotics. *Autonomous Robots*, 3(4), 375–397.
- Dudek, G., & Roy, N. (1997). Multi-robot rendezvous in unknown environments. In *Proceedings of the AAAI national conference*

- workshop on online search, Providence, Rhode Island* (pp. 22–29). AAAI.
- Freda, L., Gianni, M., Pirri, F., Gawel, A., Dubé, R., Siegart, R., & Cadena, C. (2019). 3d multi-robot patrolling with a two-level coordination strategy. *Autonomous Robots*, 43(7), 1747–1779.
- Godoy, J. E., Karamouzas, I., Guy, S. J., & Gini, M. (2016). Implicit coordination in crowded multi-agent navigation. In *30th AAAI conference on artificial intelligence*.
- Greensmith, E., Bartlett, P. L., & Baxter, J. (2004). Variance reduction techniques for gradient estimates in reinforcement learning. *Journal of Machine Learning Research*, 5((Nov)), 1471–1530.
- Higuera, J. C. G., & Dudek, G. (2013). Fair subdivision of multi-robot tasks. In *2013 IEEE international conference on robotics and automation* (pp. 3014–3019). IEEE.
- Ji, M., & Egerstedt, M. (2007). Distributed coordination control of multiagent systems while preserving connectedness. *IEEE Transactions on Robotics*, 23(4), 693–703.
- Julian, B. J., Angermann, M., Schwager, M., & Rus, D. (2012). Distributed robotic sensor networks: An information-theoretic approach. *The International Journal of Robotics Research*, 31(10), 1134–1154.
- Kapoutsis, A. C., Chatzichristofis, S. A., & Kosmatopoulos, E. B. (2017). Darp: divide areas algorithm for optimal multi-robot coverage path planning. *Journal of Intelligent & Robotic Systems*, 86(3–4), 663–680.
- Kemna, S., Kangaslahti, S., Kroemer, O., & Sukhatme, G. S. (2018). Adaptive sampling: Algorithmic vs. human waypoint selection. In *International conference on robotics and automation*.
- Low, K. H., Dolan, J. M., & Khosla, P. (2008). Adaptive multi-robot wide-area exploration and mapping. In *Proceedings of the 7th international joint conference on autonomous agents and multi-agent systems* (Vol. 1, pp. 23–30). International Foundation for Autonomous Agents and Multiagent Systems.
- Lowe, R., Wu, Y. I., Tamar, A., Harb, J., Abbeel, P., & Mordatch, I. (2017). Multi-agent actor-critic for mixed cooperative-competitive environments. In *Advances in neural information processing systems* (pp. 6379–6390).
- Manjanna, S., & Dudek, G. (2017). Data-driven selective sampling for marine vehicles using multi-scale paths. In *2017 IEEE/RSJ international conference on intelligent robots and systems (IROS)* (pp. 6111–6117). IEEE.
- Manjanna, S., Van Hoof, H., & Dudek, G. (2018). Policy search on aggregated state space for active sampling. In *International symposium on experimental robotics* (pp. 211–221). Springer.
- Meghjani, M., Manjanna, S., & Dudek, G. (2016). Multi-target rendezvous search. In *2016 IEEE/RSJ international conference on intelligent robots and systems (IROS)* (pp. 2596–2603). IEEE.
- Salam, T., & Hsieh, M. A. (2019). Adaptive sampling and reduced-order modeling of dynamic processes by robot teams. *IEEE Robotics and Automation Letters*, 4(2), 477–484.
- Savitzky, A., & Golay, M. J. (1964). Smoothing and differentiation of data by simplified least squares procedures. *Analytical Chemistry*, 36(8), 1627–1639.
- Sheng, W., Yang, Q., Tan, J., & Xi, N. (2006). Distributed multi-robot coordination in area exploration. *Robotics and Autonomous Systems*, 54(12), 945–955.
- Shkurti, F., Xu, A., Meghjani, M., Gamboa Higuera, J., Girdhar, Y., Giguere, P., Dey, B., Li, J., Kalmbach, A., Prahacs, C., Turgeon, K., Rekleitis, I., & Dudek, G. (2012). Multi-domain monitoring of marine environments using a heterogeneous robot team. In *IEEE/RSJ international conference on intelligent robots and systems* (pp. 1747–1753). Algarve, Portugal.
- Singh, A., Nowak, R., & Ramanathan, P. (2006). Active learning for adaptive mobile sensing networks. In *Proc. int. conf. information processing in sensor networks* (pp. 60–68). ACM.
- Sipahioglu, A., Kirlik, G., Parlaktuna, O., & Yazici, A. (2010). Energy constrained multi-robot sensor-based coverage path planning using capacitated arc routing approach. *Robotics and Autonomous Systems*, 58(5), 529–538.
- Spanos, D. P., & Murray, R. M. (2005). Motion planning with wireless network constraints. In *Proceedings of the 2005, American Control Conference* (pp. 87–92). IEEE.
- Sutton, R. S., McAllester, D. A., Singh, S. P., & Mansour, Y. (2000). Policy gradient methods for reinforcement learning with function approximation. In *Advances in neural information processing systems* (pp. 1057–1063).
- Tola, E., Lepetit, V., & Fua, P. (2009). Daisy: An efficient dense descriptor applied to wide-baseline stereo. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 32(5), 815–830.
- Vaughan, R. T., Stoy, K., Sukhatme, G. S., & Mataric, M. J. (2002). Lost: Localization-space trails for robot teams. *IEEE Transactions on Robotics and Automation*, 18(5), 796–812.
- Venkataramani, R., & Bresler, Y. (2000). Perfect reconstruction formulas and bounds on aliasing error in sub-Nyquist nonuniform sampling of multiband signals. *IEEE Transactions on Information Theory*, 46(6), 2173–2183.
- VK5BR, L.B.(1987) *Underwater radio communication*. Originally published in Amateur Radio.
- Williams, R. J. (1992). Simple statistical gradient-following algorithms for connectionist reinforcement learning. *Machine Learning*, 8(3–4), 229–256.
- Xu, A., Viriyasuthee, C., & Rekleitis, I. (2011). Optimal complete terrain coverage using an unmanned aerial vehicle. In *2011 IEEE international conference on robotics and automation* (pp. 2513–2519). IEEE.
- Yan, Z., Jouandeau, N., & Cherif, A. A. (2013). A survey and analysis of multi-robot coordination. *International Journal of Advanced Robotic Systems*, 10(12), 399.
- Zavlanos, M. M., & Pappas, G. J. (2008). Distributed connectivity control of mobile networks. *IEEE Transactions on Robotics*, 24(6), 1416–1428.

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Sandeep Manjanna is a James McDonnell Postdoctoral Fellow at GRASP Labs in Computer and Information Science at University of Pennsylvania, Philadelphia, PA, USA. His research interests include robotic active sampling, multi-robot coordination, robotic sensor networks, reinforcement learning, and machine learning. is a James McDonnell Postdoctoral Fellow at GRASP Labs in Computer and Information Science at University of Pennsylvania, Philadelphia, PA, USA. His research interests include robotic active sampling, multi-robot coordination, robotic sensor networks, reinforcement learning, and machine learning. Sandeep recieved his MSc and PhD from McGill University advised by Gregory Dudek. His doctoral research focuses on designing algorithms for autonomous vehicles to sample, understand, and map challenging environments.

M. Ani Hsieh is a Research Associate Professor in the Department of Mechanical Engineering and Applied Mechanics at the University of Pennsylvania. She is the Deputy Director of GRASP Lab. She received a B.S. in Engineering and B.A. in Economics from Swarthmore College in 1999 and a Ph.D. in Mechanical Engineering from the University of Pennsylvania in 2007. Ani is interested in problems that lie at the intersection of robotics, dynamical systems theory, and control. Her work focuses on how robots interact with their environments and how the environment can be leveraged as an energy reservoir for robots. Ani and her group create new abstractions to describe dynamic environments and design new architectures and algorithms for teams of robots to interact with each other and their environments. Her work is very interdisciplinary with a strong emphasis on experimental validation.

Gregory Dudek is a James McGill Chaired Professor and former Director of the McGill School of Computer Science and McGill Research Center for Intelligent Machines. He is the head of the Mobile Robotics Laboratory, where a dozen enthusiastic graduate students currently work on sensorbased robotics, multi-robot systems, and vision for robots. This involves a substantial body of work on vision for underwater systems, both with respect to vehicle design, control and perception. His group tries to answer questions such as how a robot can describe its world, how to pick actions that lead to a better understanding of the world, and how a human and a robot can communicate effectively. Prof. Dudek's work and expertise are internationally recognized with more than 200 papers published in top-tier conferences, a great many invited lectures at renowned universities, and three published books, including "Computational Principles for Mobile Robotics", together with Prof. Michael Jenkin. He is also the recipient of numerous prestigious awards, including the J.-A. Bombardier Prize from l'Association francophone pour le savoir. He is currently the President of the Canadian Image Processing and Pattern Recognition Society (CIPPRS).