

Adapting the Exploration–Exploitation Balance in Heterogeneous Swarms: Tracking Evasive Targets

Hian Lee Kwa

Singapore University of
Technology and Design
Thales Solutions Asia, Singapore
hianlee_kwa@mymail.sutd.edu.sg

Victor Babineau

University of Ottawa

Julien Philippot

University of Ottawa

Roland Bouffanais*

University of Ottawa
Roland.Bouffanais@uottawa.ca

Abstract There has been growing interest in the use of multi-robot systems in various tasks and scenarios. The main attractiveness of such systems is their flexibility, robustness, and scalability. An often overlooked yet promising feature is system modularity, which offers the possibility of harnessing agent specialization, while also enabling system-level upgrades. However, altering the agents' capacities can change the exploration–exploitation balance required to maximize the system's performance. Here, we study the effect of a swarm's heterogeneity on its exploration–exploitation balance while tracking multiple fast-moving evasive targets under the *cooperative multi-robot observation of multiple moving targets* framework. To this end, we use a decentralized search and tracking strategy with adjustable levels of exploration and exploitation. By indirectly tuning the balance, we first confirm the presence of an optimal balance between these two key competing actions. Next, by substituting slower moving agents with faster ones, we show that the system exhibits a performance improvement without any modifications to the original strategy. In addition, owing to the additional amount of exploitation carried out by the faster agents, we demonstrate that a heterogeneous system's performance can be further improved by reducing an agent's level of connectivity, to favor the conduct of exploratory actions. Furthermore, in studying the influence of the density of swarming agents, we show that the addition of faster agents can counterbalance a reduction in the overall number of agents while maintaining the level of tracking performance. Finally, we explore the challenges of using differentiated strategies to take advantage of the heterogeneous nature of the swarm.

Keywords

Exploration, exploitation, heterogeneous swarms, multi-robot Systems, swarm robotics, target tracking

1 Introduction

With the increasing popularity of multi-robot systems (MRSs), so too is the interest in implementing swarm intelligence algorithms, paving the way for swarming MRSs. These systems have found multiple applications including area mapping (Kit et al., 2019), dynamic area covering (Zoss et al., 2018), target tracking (Esterle & Lewis, 2020), and performing tasks in areas that pose environmental hazards to human operators (Schranz et al., 2020). Compared to the centralized MRS, there are three main advantages in using a decentralized swarming MRS: (1) flexibility—the

* Corresponding author.

ability for the system to swiftly adapt to rapidly changing environments, (2) robustness—the ability to cope with component failures within the system, and (3) scalability—the ability to carry out tasks in systems comprised of different numbers of agents (Dorigo et al., 2021). These attributes stem from the lack of a central controller that dictates the actions of the individual agents. The decentralized nature of such systems divides up the overall control of the MRS and gives individual robots the ability to make decisions based on their local environment. This allows them to react quickly to changes in their surroundings (flexibility), and it removes a central point of control, allows the system to continue operations despite agent failures (robustness), and eliminates potential processing bottlenecks (scalability).

Besides these three key features, an often overlooked attribute of a swarming MRS is its *modularity*. This allows for agents with different properties and behaviors to be implemented within the same collective, resulting in a heterogeneous system. Heterogeneity in swarm compositions can be utilized in several ways, including employing agents with different communication ranges (Christensen et al., 2015), or improving a system’s mapping capabilities by partitioning a search-space and assigning the partitioned regions to robots according to their movement and sensing capabilities (Shi et al., 2020). In addition, a system’s modularity presents the opportunity to incrementally introduce upgrades to a fraction of agents without the need for a system-wide upgrade. Such incremental upgrades will allow for the testing of new agent functionalities at a reduced cost (Vallegra et al., 2018). With the presence of upgraded agents, such heterogeneous swarms have the potential of delivering the same performance level with a lower agent count, thereby yielding a more cost-effective solution.

Although it has previously been shown that simply adding upgraded agents to a swarm tracking a non-evasive target can result in improved tracking performances and reduced target capture times in a target search-and-track task (Iwama & Sato, 2012; Kwa, Tokić, et al., 2020), this might not be the case in general. It has been shown in a wide variety of situations that changing the capabilities of agents also changes the ideal exploration–exploitation balance associated with optimal system performance (Kwa et al., 2022). As such, when dealing with heterogeneous swarms, modifications may need to be made to the original swarm behavioral rules to account for the different exploration–exploitation balance and better take advantage of the upgraded agents. However, it is worth stressing that heterogeneity does not automatically yield an improved system’s performance; not carrying out such modifications in the behavioral rules may even result in degraded system performances. For instance, in an area exploration task, it was shown that the introduction of robots with a larger sensor range resulted in higher amounts of agent clustering due to the higher amount of exploitation carried out, thereby increasing the amount of time required to explore the search space (Altshuler, 2009). When determining the ideal composition of heterogeneous agents of a swarm carrying out various tasks, it was found that in some tasks, such as area coverage and foraging, the presence of different agent types resulted in poorer system performances (Nishikawa et al., 2016).

One of the tools used to adjust the exploration–exploitation balance is tailoring an agent’s level of communication or interaction with its neighbors (Kwa et al., 2022). In what follows, the term *communication* has to be understood in its broadest sense, namely *information exchange* and is in no way limited to a particular communication channel or protocol. In general, high amounts of inter-agent communication lead to higher levels of exploitation while low levels of communication yield high levels of exploration. Despite the large body of research dedicated to analyzing the effect of varying an agent’s communications range on the swarm’s exploration–exploitation balance and overall behavior (Coquet et al., 2019; Talamali et al., 2021), little is known about the effect of adding one or more neighbors to an agent’s communications network (Kwa et al., 2022). Among these works, it has been demonstrated that to attain multi-agent consensus about a static noisy environment, there exists an optimum number of neighbors—and therefore an optimal exploration–exploitation balance—at which the error between the attained consensus and the ground truth is minimized (Crosscombe & Lawry, 2022). Similar observations have also been made in dynamic environments, where the optimum level of connectivity is related to the speed of the driving signal (Kwa et al., 2021; Mateo et al., 2019).

In this work, we study how heterogeneity in a swarm can be leveraged to offer adaptivity in the exploration–exploitation balance. The ultimate goal is, of course, to use this available adaptivity to maximize a system’s performance when dealing with fast-evolving circumstances. Specifically we consider a swarm attempting to search for and track multiple mobile targets. This task is carried out within the formalized framework of *cooperative multi-robot observation of multiple moving targets* (CMOMMT). Under this framework, agents position themselves in an attempt to maximize the total amount of time that each target is observed by at least one agent (Parker & Emmons, 1997). Most of the attention given to this problem is focused on maximizing the number of tracked targets when (i) the latter exceeds that of the tracking agents, and (ii) when the targets travel slower than their pursuing agents (Khan et al., 2018). Due to the assumption that fast-moving targets will always be able to evade their pursuers, cases where targets are able to travel faster and have the ability to outrun their pursuers tend to be overlooked or simply ignored. This is despite it being shown to be possible when the pursuing agents have vision of the entire environment (Janosov et al., 2017; Zhang et al., 2019).

Given the relative lack of attention dedicated to this challenging problem of tracking multiple fast-moving evasive targets, our starting point is the strategy previously proposed in Kwa et al. (2021), which was designed for use in a homogeneous swarm. This approach is a combination of (i) an adaptive inter-agent repulsive behavior, with (ii) an agent-level persistent point of attraction, and facilitated by (iii) agent-based memory. Using this strategy, the balance in the exploration of the environment and the exploitation of target information by the heterogeneous swarm can be tuned by adjusting the memory length, as well as the degree of connectivity, k , of the inter-agent communications network. Tailoring the length of memory (Falcón-Cortés et al., 2019; Nauta, Khaluf, & Simoens, 2020) and the number of neighbors in an agent’s network (Blackwell & Kennedy, 2019; Kwa, Kit, et al., 2020) have previously been established as a method that can be used to tune this exploration–exploitation balance in various multi-agent systems. Specifically, increasing memory lengths and connectivity levels serve to increase exploitative actions while decreasing memory lengths and connectivity yields higher levels of exploration. Applying such techniques in this work allows us to uncover the existence of an optimal balance between exploration and exploitation, which maximizes the system’s tracking performance.

In this article, we expand on our previous work done in Kwa et al. (2021) by focusing on the challenges and opportunities offered by a heterogeneous swarming strategy for the tracking of fast-moving evasive targets. Specifically, we investigate how the substitution of regular agents by fast-moving ones affects the swarm’s exploration–exploitation balance and tracking performance. To accomplish this, we implement the aforementioned strategy and compare the performances of simulated homogeneous and heterogeneous swarms comprised of up to 50 agents while tracking up to 3 fast-moving evasive targets. The heterogeneous system is comprised of agents with two different maximum speeds. It should be emphasized that the evasive targets have a higher maximum velocity than any of the agents, including the fast-moving ones. Through our simulations of an actual robotic system, we show that altering an agent’s memory length and level of connectivity confirms the presence of an optimal balance between exploratory and exploitative actions carried out by the system. Therefore, operating at this optimum maximizes the group-level tracking performance. It should be noted that the optimum engagement ratio changes depending on the task presented to the system, as well as the heterogeneous composition of the system (Kwa, Tokić, et al., 2020). In addition, we establish that smaller heterogeneous swarms are able to match the performance of a full-size homogeneous one over a range of different densities of agents. Finally, we explore the challenges of using differentiated tracking strategies in heterogeneous swarms.

2 Problem Formulation and Multi-Agent Dynamics

2.1 Problem Statement

Here, we define the problem under the CMOMMT framework. We consider a set of N tracking agents $\mathcal{A} = \{a_1, a_2, \dots, a_N\}$ and a set of J targets $\mathcal{O} = \{o_1, o_2, \dots, o_J\}$ moving within a bounded

two-dimensional square search space of dimensions $L \times L$. The agents can either be slower agents, or faster upgraded agents. Each agent is also given a memory of length, t_{mem} , and communicates with a fixed number of neighbors, $k \in [1, N - 1]$. These aspects are discussed in the following sections.

2.2 Search and Track Strategy

The search and track strategy implemented in this work was first introduced in Kwa et al. (2021) for use in homogeneous swarms tracking both evasive and non-evasive targets. This strategy consists of two components, namely *attraction* and *repulsion* (or avoidance). The approach involves an attraction component that is closely tied to the topology of the interaction network, which effectively controls the flow of information throughout the system. The repulsion component is designed as an adaptive inter-agent repulsive behavior that yields simultaneous exploratory actions while the units are exploiting information about a tracked target. An appropriate integration of these two components promotes system exploration and exploitation respectively by generating a velocity vector by combining them at the start of each time-step:

$$\mathbf{v}_i[t] = \mathbf{v}_{i,\text{att}}[t] + \mathbf{v}_{i,\text{rep}}[t], \quad (1)$$

where $\mathbf{v}_{i,\text{att}}[t]$ (resp. $\mathbf{v}_{i,\text{rep}}[t]$) is the velocity vector obtained from the attraction (resp. repulsion) component. In the simulation, the resultant velocity vector, $\mathbf{v}_i[t]$ is then scaled by v_{max} , the maximum velocity of an agent, before being used to calculate the agent's position in the subsequent time-step. This strategy employed by the swarm, which is detailed in Algorithm 1, makes use of short-term agent memory and a variable k -nearest neighbor communications network to adjust the system's level of exploration and exploitation. These two parameters will be discussed in the subsequent sections. For further details regarding the individual strategy components, the reader is directed to Kwa et al. (2021).

2.3 Swarm Communication Network

The agents within the swarming system are connected to each other through the use of a dynamic k -nearest neighbor network. Changing the connectivity by means of the network degree, k , has been shown to have large impacts on the overall collective response of a swarm subjected to dynamic changes of varying time-scales (Mateo et al., 2019). It has been demonstrated that higher levels of connectivity cause swarming agents to cluster in small areas, signifying high levels of exploitation, while lower levels of connectivity are associated with more exploratory actions (Altshuler, 2009; Kwa, Kit, et al., 2020). As such, the overall exploration–exploitation balance of the swarm presented in this article is primarily adjusted by changing the value of k . It is important to note that in our framework, a neighborhood is understood in the network sense, hence, an agent i has

Algorithm 1. Dynamic k -nearest network search and tracking strategy.

- 1: Set $t = 0$, $k \in [2, N - 1]$, $\omega = 1$, and $c = 0.5$
 - 2: **while** System active **do**
 - 3: **for** All agents $i \in [1, N]$ **do**
 - 4: Determine point of attraction
 - 5: Calculate $\mathbf{v}_{\text{att},i}$
 - 6: Calculate $\mathbf{v}_{\text{rep},i}$
 - 7: $\mathbf{v}_i[t] \leftarrow \mathbf{v}_{\text{att},i}[t] + \mathbf{v}_{\text{rep},i}[t]$ // Apply Eq 1
 - 8: $\mathbf{v}_i[t] \leftarrow (v_{\text{max}}/v_i[t]) \cdot \mathbf{v}_i[t]$ // Scale velocity vector with maximum agent speed
 - 9: $\mathbf{x}_i[t + 1] \leftarrow \mathbf{x}_i[t] + \mathbf{v}_i[t]$ // Update agent position
 - 10: **end for**
 - 11: $t \leftarrow t + 1$
 - 12: **end while**
-

as many neighbors as its degree, k . Note that this network is a temporal network owing to the motion of neighboring agents. Instead of modifying the collective behavior of the MRS by changing the communication range of the swarm, we designate the number of neighbors in an agent's communication network. This allows for a higher level of control over the system's exploration–exploitation balance. In practice, this can be attained by adjusting the communications range of an individual agent until the desired number of neighbors is attained (Rausch et al., 2019) or by only processing information originating from the neighboring units according to the selected value k of the interconnecting topological network (Kwa, Kit, et al., 2020).

2.4 Short-Term Agent-Based Memory

In Kwa et al. (2021), we introduced a short-term agent-based memory into the search and track strategy. This was inspired by the work done in Nauta, Khaluf, & Simoens (2020) and Nauta, Simoens, & Khaluf (2020), which demonstrated increases in system performance when a swarm was placed in a nondestructive foraging scenario. This was especially so when resources were clustered together in the environment, allowing agents to return to resource dense patches more often. While tracking non-evasive targets, it was determined that the implementation of memory causes the swarm to aggregate in a location where the target is no longer present, thereby degrading the overall swarm's tracking performance (Coquet et al., 2019, 2021; Kwa, Kit, et al., 2020). However, in Kwa et al. (2021), we demonstrated that the inclusion of memory was essential to tracking a fast-moving evasive target. Since an evasive target moves to avoid contact with the pursuing agents, the inclusion of memory in this scenario allowed for a more persistent point of attraction to be generated, effectively increasing the number of exploitative actions performed by the swarm. This in turn allows the system to close in on a target even though none of the agents is able to detect the presence of a target. This can happen as long as an agent has a valid point of attraction, thereby enabling the tracking of the target by the swarm.

3 Simulation Conditions

3.1 Swarm Robotic Platform

In this work, we model the *bunch of buoys* (BoB) system, originally intended for the dynamic monitoring of open water bodies (Vallegra et al., 2018; Zoss et al., 2018). The current system employs two types of buoys: BoB-0, which weighs 7.4 kg and can attain a maximum speed of 1.0 m/s, and BoB-1, weighing 3.2 kg and capable of attaining a top speed of 2.6 m/s. Being faster and lighter, BoB-1 is able to better react to dynamic changes in its operating environment compared to BoB-0. Both robotic platforms have propulsion systems allowing for omnidirectional movement. In addition, the platforms house a suite of sensors, enabling the characterization of their local environment. A distributed mesh communications system is used for sharing locally sensed information among the buoys, as well as for the sending and receiving of commands. In the simulations, 50 buoys tracking a fictitious moving target were modeled. Swarms with various compositions of BoB-0 and BoB-1 units and levels of connectivity were tested to observe the effects of gradually adding upgraded units to the system.

3.2 Simulation Parameters and Performance Metric

The bounded two-dimensional square search space of dimensions $L \times L$ is such that $4 \leq L \leq 445$. The slow and fast agents have associated maximum velocities of $\mathbf{v}_{a_0, \max} = 0.1$, $\mathbf{v}_{a_1, \max} = 0.26$ respectively and x and y positions represented by $\mathbf{x}_i = (x_i, y_i)$. The agent's memory of length is given as $t_{\text{mem}} \in [0, 50]$. In simulations involving differentiated communications strategies, k_f (resp. k_s) represents the number of topological neighbors for fast (resp. slow) agents.

The targets are modeled using disc-shaped binary objective functions with fixed radii of $\rho = 1$. These targets also have $x[t]$ and $y[t]$ positions associated with the center of the disc and a maximum velocity, $\mathbf{v}_{o, \max}$, and do not overlap. This maximum velocity is greater than that of the two types of

agents; therefore: $\mathbf{v}_{a_0,\max} < \mathbf{v}_{a_1,\max} < \mathbf{v}_{o,\max}$. A target is considered to be tracked if there exists an agent within its radius. Formally:

$$\text{cov}(\theta_m, t) = \begin{cases} 1 & \exists i \in \mathcal{A} \text{ s.t. } \|\mathbf{x}_i - \mathbf{x}_m\| \leq \rho, \\ 0 & \text{otherwise.} \end{cases} \quad (2)$$

The modeling of the target as a binary objective function is done to make the task even more challenging and is similar to a visual search task where a target needs to be seen to have its presence confirmed (Esterle & Lewis, 2020). This is in contrast to tracking a target from the intensity of an emitted signal (e.g., radio signal strength, chemical plume), in which various gradient-descent methods can be used. Indeed, such gradient-descent methods form one of the most widely used techniques for tracking targets, and become completely ineffective when dealing with such binary objective functions. Using a binary objective function, as is done here, ensures that such simplistic techniques are not used by the system. It can be argued that the full power of swarm intelligence can only be accessed when dealing with such challenging scenarios.

The targets are set to move according to an evasive movement policy. Specifically, a target initially travels towards randomly generated waypoints within the search space until it encounters a swarming agent within its radius. Upon this encounter, the target calculates its velocity using the repulsion Equation 3:

$$\mathbf{v}_{\text{rep},m}[t] = - \sum_{i \in \mathcal{N}_i} \left(\frac{a_R}{r_{mi}[t]} \right)^d \frac{\mathbf{r}_{mi}[t]}{r_{mi}[t]}, \quad (3)$$

where $\mathbf{r}_{mj}[t]$ is the vector from target m to agent i at time-step t , with agent i being a part of \mathcal{N}_i , the group of agents located within the target's radius. The magnitude of the repulsion vector is determined by $r_{mi}[t]$, where $r_{mi}[t] = \|\mathbf{r}_{mi}[t]\|$, and a_R , the repulsion strength of the target.

Should the target encounter agents within its radius for t_{limit} consecutive iterations, the target travels in a straight line for t_{evade} time-steps in an attempt to outrun its pursuers.

The goal of the system is to maximize the tracking performance of the targets within this environment, which can be measured quantitatively by means of the normalized reward function:

$$\Xi = \frac{1}{TJ} \sum_{t=1}^T \sum_{m=1}^J \text{cov}(\theta_m, t), \quad (4)$$

where T is the total time period of interest and J is the total number of targets within the search space.

3.3 Exploration and Exploitation Balance

In addition to group-level tracking performance, the overall complex collective dynamics of a swarm can be studied through the quantification of the exploration–exploitation balance of the system. This metric was first introduced in Kwa et al. (2021). Here, the proportion of a swarm's agents engaged with a target, which shall be referred to as the swarm's *engagement ratio*, is used to quantify the balance between exploratory and exploitative actions carried out by the system. An agent is considered to be engaged with a target if it has entered the *tracking* state and $S_i(\mathbf{x}_i[t], t) = 1$, thereby giving the overall normalized engagement ratio:

$$\Theta = \frac{1}{NT} \sum_{t=1}^T \sum_{i=1}^N S_i[t]. \quad (5)$$

When using this metric, a system that displays a high engagement ratio indicates that agents are spending a high fraction of time attempting to track a target, i.e., exploiting it. Conversely, a low engagement ratio indicates that agents are spending less time attempting to track a target, which is characteristic of higher levels of exploration. It should be noted that while our initial work in Kwa et al. (2021) pointed to the existence of an optimum engagement ratio to maximize tracking performance, this optimum changes according to the task assigned. As such, it would not be possible to declare what an optimum level of engagement should be. This changing optimum engagement ratio will be addressed in the subsequent sections.

4 Simulation Results

Swarms of $N = 50$ agents comprised of various proportions of fast and slow agents are deployed within the square search space. The width of the search space is kept constant at $L = 30$ for all simulations except for those studying swarm density, where $4 \leq L \leq 445$. The agents are tasked with tracking $J \in \{1, 2, 3\}$ targets following the evasive movement policy that travel at a maximum speed of $\mathbf{v}_{o,\max} = 0.3$ non-dimensional units per time-step. All simulations are carried out over a period of 400,000 iterations to ensure the statistical stationarity of the results.

4.1 Impact From the Level of Network Connectivity

Figure 1 illustrates that there exists an optimum level of connectivity, κ^* , at which the tracking performance of the system is maximized. This is true for both homogeneous and heterogeneous swarms. On one hand, when $\kappa < \kappa^*$, the low tracking by the system is due to the low level of exploitation, as attested by the low values of Θ , resulting in the agents being unable to cluster around the target adequately. On the other hand, when $\kappa > \kappa^*$, too much exploitation carried out by the system yields excessive clustering of agents around the target, ultimately impairing the tracking performance. Indeed, when operating at high levels of κ , the agents take a longer time to re-explore the environment and reacquire a target's position after being outrun. The figure also suggests that this optimum κ is reduced when fast agents are introduced to the swarm. This implies that while no modifications to the overall approach are necessary to improve the swarm's performance when introducing fast agents, doing so by altering an agent's connectivity degree will maximize the performance of the heterogeneous swarm. The reasons for this will also be further explored in section 4.3.




Figure 1. Tracking performance (right) and engagement ratio (left) of a heterogeneous swarm comprised of 45 fast agents and 5 slow agents (solid line), as well as a homogeneous swarm comprised of 50 slow agents (dashed line) tracking multiple fast-moving evasive targets. Simulations were performed using various levels of network connectivity k and a memory length of $t_{\text{mem}} = 20$.

Secondly, Figure 1 reveals that while tracking evasive targets, the strategy considered in homogeneous swarms can readily be used with a heterogeneous one to a greater effect without the need for adjustments; an improvement in the MRS’s tracking performance occurs regardless of the system’s level of connectivity. This result is congruent with our previous findings that no specific modifications need to be made to benefit from a heterogeneous swarm in a dynamic area coverage scenario (Vallegra et al., 2018). This is despite the fact that other studies have shown that either modifications to the overall strategy, such as changes to the communications network, need to be made, or that specific initialization conditions need to be satisfied to fully benefit from a heterogeneous system (Altshuler, 2009; Nishikawa et al., 2016). The reasons for the improvement in tracking performance in spite of the change in exploration–exploitation balance are explained in section 4.3.

Thirdly, it is demonstrated that the tracking performance of the swarm decreases when more targets are added to the environment. Again, this occurs regardless of the system’s level of connectivity. This is to be expected as simultaneously tracking multiple targets is more challenging than just tracking a single one. It can also be observed that both increasing the system’s degree of connectivity and increasing the number of targets also serves to raise the engagement ratio, Θ . The former occurs because elevating the system’s level of connectivity results in a target’s positional information being spread around the system’s agents more readily. The latter is due to the higher number of targets naturally increasing the likelihood of a tracking agent encountering a target. Consequently, in these two situations, the likelihood of an agent attempting to engage with a target also increases, either due to the increased availability of target information or the higher probability of an agent encountering a target, thereby increasing Θ . It should be noted that Θ does not differentiate between targets.

Interestingly, the optimum exploration–exploitation balance appears to remain constant regardless of the number of targets—rightmost point in Figure 2. Therefore, to maximize the tracking performance of the swarm, more exploration needs to be carried out. As more exploration is carried out at lower levels of connectivity, a swarm’s k^* reduces when it attempts to track multiple targets, which can be seen in Figure 1.

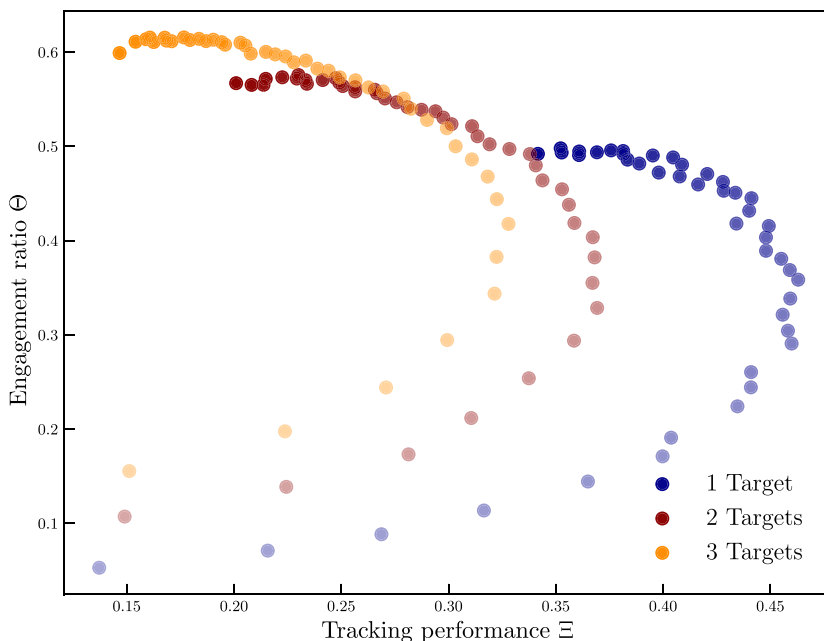


Figure 2. Engagement-tracking plots of a heterogeneous swarm comprised of 15 fast agents tracking multiple fast-moving evasive targets. Darker shaded points indicate swarms communicating through networks with higher values of k .

4.2 Impact From Agent-Based Memory

As previously mentioned, the implementation of an agent-based memory was done to create a more lasting point of attraction, thus allowing the agents to flock around an evasive target. Since it facilitates the clustering of agents around the target, it also functions as another tool that permits the adjustment of the exploration–exploitation balance of the swarm. However, caution needs to be taken when implementing agent memory as its use was previously determined to be detrimental when tracking non-evasive targets due to the system’s use of information that is no longer valid when excessive memory lengths were used.

Despite its associated disadvantages, it can clearly be seen in Figure 3 that the use of a short-term memory is essential to the swarm’s ability to track an evasive target. Without the use of memory ($t_{mem} = 0$), the system is unable to effectively track an evasive target, and the addition of small amounts of memory is associated with a sharp increase in both tracking performance and engagement. This is because allowing the agents to exploit information from its memory generates a persistent point of attraction based on a target’s last known position. Without the use of memory, even though the swarm may periodically encounter a target, the swarm tends to expand until the agents reach their static equilibrium positions and is unable to close in on the target because of its quick evasive maneuvers.

However, should the agent be allowed to exploit information from a longer-than-optimal memory length, there is an overall degradation of the tracking performance as the system starts to suffer from the effects stemming from the exploitation of information that is no longer valid. This is despite the fact that the engagement ratio of the system continues to rise with the increasing memory lengths. This rise in Θ is due to the increased amount of time the swarming agents spend in the *tracking state* (see section 3.3), causing them to aggregate at a position where a target is no longer present. This further emphasizes the importance of maintaining an optimal balance between exploration and exploitation as mentioned in section 4.1.

4.3 Impact From the Introduction of Fast Agents

As can be seen in Figure 4, increasing the proportion of fast agents results in an increase in both Ξ and Θ . It is worth highlighting again that these increases are obtained without modifying the original tracking strategy. Observations made from the simulations reveal the reasons why such

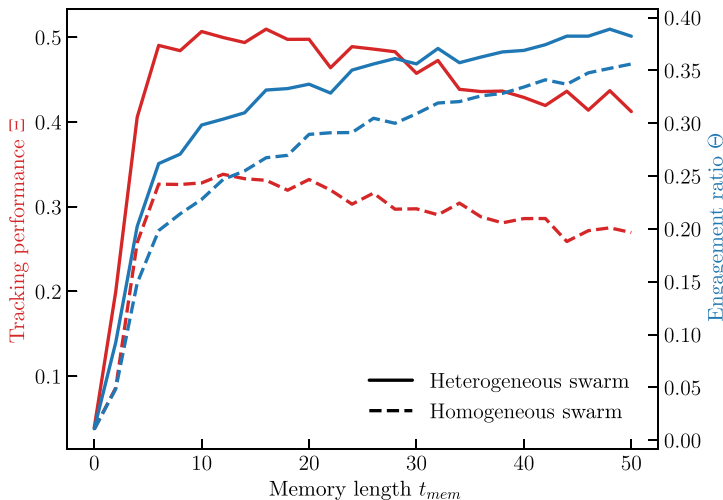


Figure 3. Tracking performance (red line) and engagement ratio (blue line) of a heterogeneous swarm comprised of 15 BoB-1 units and 35 BoB-0 units (solid line) and a homogeneous swarm comprised of slow agents (dashed line) tracking a single fast-moving evasive target. Simulations were performed using a network connectivity of $k = 14$ and various memory lengths.




Figure 4. Tracking performance (right) and engagement ratio (left) of swarms comprised of 50 agents with various proportions of BoB-I agents tracking a fast-moving evasive target. Simulations were performed using a network connectivity of $k = 12$ and a memory length of $t_{mem} = 20$.




Figure 5. Heterogeneous swarm comprised of 35 BoB-0 units (black crosses) and 15 BoB-I units (blue crosses) tracking an evasive target (dark blue disc). (Left) BoB-I units clustering while pursuing target; (Right) BoB-I units redistributed around the environment due to inter-agent repulsion.

modifications are unnecessary. As illustrated in Figure 5, while the faster agents may flock together when actively trying to engage with the target, they are eventually able to disperse and distribute themselves across the swarm after losing track of the target. This is thanks to the adaptive inter-agent repulsion component in the search and track strategy, which increases the chances of a fast agent encountering a target.

The observations and the quantitative results suggest that the increases in performance and engagement are solely due to the increased mobility of the faster agents within the swarm. Since the faster agents are better able to keep up with the targets, they directly increase the swarm’s tracking capability. They are also able to disseminate more information about the targets’ locations across the swarm, allowing neighboring agents to remain engaged with a target for longer and close in on it, further improving the tracking performance. The increase in Θ also means that the addition of

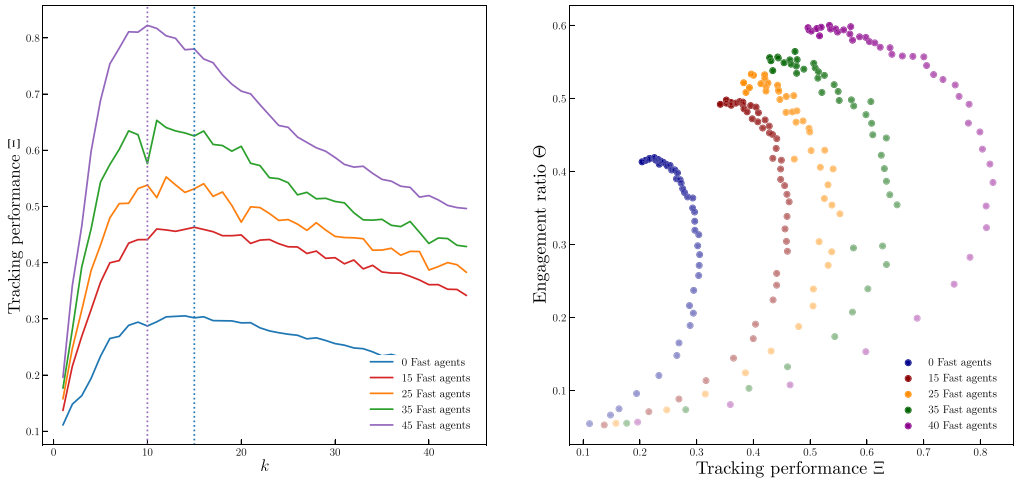


Figure 6. Swarms comprised of 50 agents with different agent compositions of BoB-1 and BoB-0 units tracking a single fast-moving evasive target. A system with 0 fast agents is a homogeneous swarm comprised solely of BoB-0 units. (Left) Tracking performance of the system while operating with different levels of connectivity k . (Right) Engagement-Tracking plots where darker shaded points indicate swarms communicating through networks with higher values of k .

faster agents biases the system’s actions towards higher levels of exploitation. However, due to the increased mobility of the system in general, the optimum exploration–exploitation balance of the swarm also moves in favor of exploitation. This is illustrated in Figure 6 where gradually increasing the number of fast agents within the system results in an increased optimum engagement ratio.

Despite this apparent shift in the ideal exploration–exploitation balance in favor of exploitation, Figure 6 also shows that the optimum connectivity degree, k^* , appears to decrease with the introduction of fast-moving BoB-1 units into the system—from $k^* \sim 15$ when a homogeneous swarm is used to $k^* \sim 10$ when the swarm consists of 45 BoB-1 units and 5 BoB-0 units. This reduction in k^* is characteristic of a swarm trying to increase its level of exploration. This observation is similar to the results obtained when tracking non-evasive targets (Kwa, Tokić, et al., 2020), albeit to a lesser extent. This suggests that the primary effect of substituting fast agents into a swarm is the increase of both the system’s capability and propensity to exploit target information. To balance the increased amount of exploitation promoted by the faster agents, the system needs to carry out more exploration. As more exploration is carried out at lower levels of connectivity, the optimum connectivity degree, k^* , decreases when faster agents are substituted into the system.

4.4 Impact of Swarm Density

As alluded to previously, the inclusion of highly motile agents within the swarm can also reduce the overall number of agents required in the system. However, such a reduction in agent numbers will result in the swarm operating with a reduced swarm density. As shown in Figure 7, when varying the density of agents, three different phases can be observed. At low density ($< 3 \times 10^{-2}$ agents / unit area), the swarm is unable to effectively track the target. In this phase, the swarm density is too low and agents in the system are unable to collectively follow the target. At high density (> 1 agent / unit area), the system is able to constantly track the target as it moves around the search environment. In this phase, the search area is too small and the target will encounter an agent regardless of where it moves. Between these two phases is a *transition* region where the tracking performance sharply increases with density. The figure also reveals that within this transition region, heterogeneous swarms that include faster moving agents are able to match, if not exceed, the tracking ability of a homogeneous swarm solely comprised of non-upgraded agents in all three density phases. This is because the faster agent speeds are able to compensate for the reduced agent numbers. Since the




Figure 7. Evasive tracking performance (left) and engagement ratio (right) of a homogeneous swarm comprised of 50 agents and two heterogeneous swarms comprised of a total of 40 and 30 agents each. Simulations were carried out over various area sizes to vary swarm density with agents using a $k = 14$ network.

faster agents are able to better keep up with the target, the swarm is able to remain engaged with the target, allowing for an adequate amount of exploitation to be performed. As such, a larger swarm comprised of slower agents is able to outperform a smaller swarm that includes faster agents.

4.5 Heterogeneous Strategies

To ascertain the possibility of using a heterogeneous memory strategy, simulations are carried out using homogeneous swarms with varying agent speeds while utilizing different memory lengths. However, as shown in Figure 8, the optimum memory length does not appear to change with the




Figure 8. Evasive tracking performance of homogeneous swarms with different $v_{a,max}$ and memory lengths. Simulations were run at $k = 14$ and $v_{o,max} = 0.3$.




Figure 9. Single evasive target tracking performance of a heterogeneous system comprised of 30 slow agents and 10 fast agents, each having different levels of connectivity.

maximum speed of the agents. This therefore suggests that the use of a heterogeneous strategy involving the use of different memory lengths will not improve the tracking performance of a swarm.

We demonstrated in section 4.3 that the addition of faster moving agents reduces the optimal level of connectivity, k^* , for the swarm. In addition, previous studies have shown that this optimal level of connectivity changes with the speed of the driving stimulus (Kwa, Kit, et al., 2020; Mateo et al., 2019). As such, to explore the possibility of using a heterogeneous strategy using different levels of connection, we carried out the same evasive target tracking simulation where the level of connectivity of the fast agents (k_f) is different from the level of connectivity of the slow agents (k_s).

From Figure 9, we were able to identify an optimum combination of levels of connectivity to maximize the tracking performance. However, this optimum occurs when $k_s = k_f = 10$ and none of the other cases tested suggests that employing a differentiated strategy would further improve a swarm's tracking performance. These results hint that further attempts to improve the tracking performance of heterogeneous MRS swarms should integrate other behavioral factors, for example, differentiated task allocation based on an agent's placement, prediction of target movement patterns.

5 Discussion

The use of decentralized swarming MRSs offers users many advantages, including system modularity. By exploiting this modular character, heterogeneous swarms comprised of agents with different behaviors and capabilities can be implemented, potentially leading to improved system performance. However, the different agent abilities may need to be taken into account when designing swarming strategies due to changes in the ideal exploration–exploitation balance.

In this work, we studied the implementation of heterogeneous swarms performing the task of searching for and tracking multiple fast-moving evasive targets under the CMOMMT framework. Through simulating an actual environmental monitoring robotic system employing swarming agents,

we demonstrated that the introduction of faster agents is indeed able to improve the overall tracking performance of the swarm without modifying the originally implemented behaviors. This is because the higher mobility of the faster units gives them the ability to better keep up with and track the targets, leading to better information dissemination of the targets' locations and higher levels of engagement with the target. This also leads to an overall shift in the optimum exploration–exploitation balance of the system, favoring more exploitative behaviors to maximize the swarm's tracking performance. While the strategy does not need to be modified, the performance of the swarm can be further improved by accounting for the increased amount of exploitation brought on by the introduction of fast-moving agents. By reducing the agents' degree of connectivity, more exploratory action is promoted within the system, thereby offsetting the excessive exploitation brought on by the introduction of the fast agents. This allows for the ideal exploration–exploitation balance to be restored, thereby improving the performance of the swarm. It should be noted that while we have shown that the optimum level of connectivity, k^* , decreases with the proportion of fast agents within the swarm, a more detailed study would be required to explore how k^* varies as a function of the swarm composition.

Through varying the size of the operating environment, we have also showed that smaller heterogeneous swarms incorporating faster agents are able to match and outperform the performance of a full-size homogeneous swarm comprised at all swarm densities. This is because faster agent speeds are able to track the evasive target for longer periods of time. This compensates for the reduced agent numbers and allows the swarm to remain engaged with the target, allowing for an adequate amount of exploitation to be performed.

We have also demonstrated the impact of the level of network connectivity and the introduction of agent-based memory on the swarm's tracking performance, as well as their effects on the level of exploration and exploitation carried out by the swarm. By increasing either the level of connectivity, k , or the length of agent memory, the system biases itself towards exploitation. Conversely, reducing k or memory length causes the swarm to favor more exploratory actions. As such, there exists an optimal level of connectivity and memory length at which there is a suitable balance between exploration and exploitation carried out by the system. This balance is responsible for maximizing the tracking performance of the swarm. This optimum moves in favor of exploration when tracking more than one target to counter the higher levels of target engagement provoked by the presence of multiple targets. Beyond this analysis of the influence of k , the degree of the network topology, it has recently been shown that behavioral changes in a swarm can propagate according to simple or complex contagions depending on other network parameters such as network clustering or Kirchhoff index (Horsevad et al., 2022). Looking into the influence of these other network characteristics might reveal other parameters that influence the swarm's response when tracking targets.

It should be noted that in this work, we have not touched on an agent's communication range, which may pose an issue when operating at very low swarm density in expansive environments. In such scenarios, the lack of communication neighbors would hinder the amount of exploitation carried out by the system, undermining its ability to track a target. Similarly, it was also assumed that agents would have perfect information of the target's location when within the target radius and that there was no noise affecting an agent's sensing capabilities. As can be expected, the addition of noise would also affect the agent's ability to exploit a target's positional information. Future works on this subject should therefore focus on trying to encourage and control the level of exploitation when the range of communication is restricted or when imperfect data is provided to the agents.

Given the ability to tune the balance between exploratory and exploitative actions in a heterogeneous swarm, further explorations of this topic may offer new ways to learn collective behaviors through multi-agent reinforcement learning that could potentially take advantage of differentiated strategies for fast and slow agents (Kouzehgar et al., 2020). While this work performed a preliminary study on the use of differentiated strategies in heterogeneous swarms where agents had different levels of mobility, a system's heterogeneity can come in different forms (e.g., deploying agents with different sensing abilities, mixed aerial and ground robot swarms). Further

work should be done on the application of differentiated strategies for such systems, i.e., behavioral heterogeneity.

Acknowledgments

H. L. Kwa is funded by Thales Solutions Asia Pte. Ltd., under the Singapore Economic Development Board Industrial Postgraduate Programme (IPP).

References

- Altshuler, Y. (2009). To add with caution—decreasing a swarm robotics' efficiency by imprudently enhancing the robots' capabilities. In *ICARA 2009: Proceedings of the fourth international conference on autonomous robots and agents* (pp. 351–356). IEEE. <https://doi.org/10.1109/ICARA.2000.4804014>
- Blackwell, T., & Kennedy, J. (2019). Impact of communication topology in particle swarm optimization. *IEEE Transactions on Evolutionary Computation*, 23(4), 689–702. <https://doi.org/10.1109/TEVC.2018.2880894>
- Christensen, A. L., Oliveira, S., Postolache, O., João de Oliveira, M., Sargento, S., Santana, P., Nunes, L., Velez, F., Sebastião, P., Costa, V., Duarte, M., Gomes, J., Rodrigues, T., & Silva, F. (2015). Design of communication and control for swarms of aquatic surface drones. In *ICAAART 2015: Proceedings of the international conference on agents and artificial intelligence* (pp. 548–555). ICAART. <https://doi.org/10.5220/0005281705480555>
- Coquet, C., Arnold, A., & Bouvet, P.-J. (2021). Control of a robotic swarm formation to track a dynamic target with communication constraints: Analysis and simulation. *Applied Sciences*, 11(7), Article 3179. <https://doi.org/10.3390/app11073179>
- Coquet, C., Aubry, C., Arnold, A., & Bouvet, P.-J. (2019). A local charged particle swarm optimization to track an underwater mobile source. In *OCEANS Marseille, 2019* (Article 527). IEEE. <https://doi.org/10.1109/OCEANSE.2019.8867527>
- Crosscombe, M., & Lawry, J. (2022). The impact of network connectivity on collective learning. In F. Matsuno, S. Azuma, & M. Yamamoto (Eds.), *Distributed autonomous robotic systems. DARS 2021* (pp. 82–94). Springer. https://doi.org/10.1007/978-3-030-92790-5_7
- Dorigo, M., Theraulaz, G., & Trianni, V. (2021). Swarm robotics: Past, present, and future [Point of View]. *Proceedings of the IEEE*, 109(7), 1152–1165. <https://doi.org/10.1109/jproc.2021.3072740>
- Esterle, L., & Lewis, P. R. (2020). Distributed autonomy and trade-offs in online multi-object k -coverage. *Computational Intelligence*, 36(2), 720–742. <https://doi.org/10.1111/coin.12264>
- Falcón-Cortés, A., Boyer, D., & Ramos-Fernández, G. (2019). Collective learning from individual experiences and information transfer during group foraging. *Journal of the Royal Society Interface*, 16(151), Article 20180803. <https://doi.org/10.1098/rsif.2018.0803>, PubMed: 30958179
- Horsevad, N., Mateo, D., Kooij, R. E., Barrat, A., & Bouffanais, R. (2022). Transition from simple to complex contagion in collective decision-making. *Nature Communications*, 13(1), Article 1442. <https://doi.org/10.1038/s41467-022-28958-6>, PubMed: 35301305
- Iwama, T., & Sato, M. (2012). Group chase and escape with some fast chasers. *Physical Review E*, 86(6), Article 067102. <https://doi.org/10.1103/PhysRevE.86.067102>, PubMed: 23368082
- Janosov, M., Virágh, C., Vásárhelyi, G., & Vicsek, T. (2017). Group chasing tactics: How to catch a faster prey. *New Journal of Physics*, 19, Article 053003. <https://doi.org/10.1088/1367-2630/aa69e7>
- Khan, A., Rinner, B., & Cavallaro, A. (2018). Cooperative robots to observe moving targets: Review. *IEEE Transactions on Cybernetics*, 48(1), 187–198. <https://doi.org/10.1109/TCYB.2016.2628161>, PubMed: 27925600
- Kit, J. L., Dharmawan, A. G., Mateo, D., Foong, S., Soh, G. S., Bouffanais, R., & Wood, K. L. (2019). Decentralized multi-floor exploration by a swarm of miniature robots teaming with wall-climbing units. In *2019 International symposium on multi-robot and-agent systems (MRS)* (pp. 195–201). IEEE. <https://doi.org/10.1109/MRS.2019.8901058>
- Kouzehgar, M., Meghjani, M., & Bouffanais, R. (2020). Multi-agent reinforcement learning for dynamic ocean monitoring by a swarm of buoys. In *Global OCEANS 2020: Singapore-U.S. Gulf Coast* (Article 128). IEEE. <https://doi.org/10.1109/IEEECONF38699.2020.9389128>

- Kwa, H. L., Kit, J. L., & Bouffanais, R. (2020). Optimal swarm strategy for dynamic target search and tracking. In *AAIAS 2020: Proceedings of the 19th international conference on autonomous agents and multiagent systems* (pp. 672–680). ACM.
- Kwa, H. L., Kit, J. L., & Bouffanais, R. (2021). Tracking multiple fast targets with swarms: Interplay between social interaction and agent memory. In *ALIFE 2021: The 2021 conference on artificial life* (Article 62). MIT Press. https://doi.org/10.1162/isal_a_00376
- Kwa, H. L., Kit, J. L., & Bouffanais, R. (2022). Balancing collective exploration and exploitation in multi-agent and multi-robot systems: A review. *Frontiers in Robotics and AI*, 8, Article 771520. <https://doi.org/10.3389/frobt.2021.771520>, PubMed: 35178430
- Kwa, H. L., Tokić, G., Bouffanais, R., & Yue, D. K. P. (2020). Heterogeneous swarms for maritime dynamic target search and tracking. In *Global OCEANS 2020: Singapore-U.S. Gulf Coast* (Article 145). <https://doi.org/10.1109/IEEECONF38699.2020.9389145>
- Mateo, D., Horsevad, N., Hassani, V., Chamanbaz, M., & Bouffanais, R. (2019). Optimal network topology for effective collective response. *Science Advances*, 5(4), Article eaau0999. <https://doi.org/10.1126/sciadv.aau0999>, PubMed: 30949570
- Nauta, J., Khaluf, Y., & Simoens, P. (2020). Hybrid foraging in patchy environments using spatial memory. *Journal of the Royal Society Interface*, 17(166), Article 0026. <https://doi.org/10.1098/rsif.2020.0026>, PubMed: 32429823
- Nauta, J., Simoens, P., & Khaluf, Y. (2020). Memory induced aggregation in collective foraging. In *ANTS2020, The 13th international conference on swarm intelligence* (pp. 176–189). Springer. https://doi.org/10.1007/978-3-030-60376-2_14
- Nishikawa, N., Suzuki, R., & Arita, T. (2016). Coordination control design of heterogeneous swarm robots by means of task oriented optimization. *Artificial Life and Robotics*, 21, 57–68. <https://doi.org/10.1007/s10015-015-0255-4>
- Parker, L. E., & Emmons, B. A. (1997). Cooperative multi-robot observation of multiple moving targets. In *Proceedings of the international conference on robotics and automation* (pp. 2082–2089). IEEE. <https://doi.org/10.1109/robot.1997.619270>
- Rausch, I., Reina, A., Simoens, P., & Khaluf, Y. (2019). Coherent collective behaviour emerging from decentralised balancing of social feedback and noise. *Swarm Intelligence*, 13, 321–345. <https://doi.org/10.1007/s11721-019-00173-y>
- Schranz, M., Umlauf, M., Sende, M., & Elmenreich, W. (2020). Swarm robotic behaviors and current applications. *Frontiers in Robotics and AI*, 7, Article 36. <https://doi.org/10.3389/frobt.2020.00036>, PubMed: 33501204
- Shi, Y., Wang, N., Zheng, J., Zhang, Y., Yi, S., Luo, W., & Sycara, K. (2020). Adaptive informative sampling with environment partitioning for heterogeneous multi-robot systems. In *2020 IEEE/RSJ international conference on intelligent robots and systems (IROS)* (pp. 11718–11723). IEEE. <https://doi.org/10.1109/IROS45743.2020.9341711>
- Talamali, M. S., Saha, A., Marshall, J. A. R., & Reina, A. (2021). When less is more: Robot swarms adapt better to changes with constrained communication. *Science Robotics*, 6(56), Article abf1416. <https://doi.org/10.1126/scirobotics.abf1416>, PubMed: 34321345
- Vallegra, F., Mateo, D., Tokić, G., Bouffanais, R., & Yue, D. K. P. (2018). Gradual collective upgrade of a swarm of autonomous buoys for dynamic ocean monitoring. In *OCEANS 2018 MTS/IEEE Charleston* (Article 642). IEEE. <https://doi.org/10.1109/OCEANS.2018.8604642>
- Zhang, S., Liu, M. Y., Lei, X. K., & Huang, Y. K. (2019). Dynamics and motion patterns of a k -capture game with attraction–repulsion interaction. *Europhysics Letters*, 128(1), Article 10003. <https://doi.org/10.1209/0295-5075/128/10003>
- Zoss, B. M., Mateo, D., Kuan, Y. K., Tokić, G., Chamanbaz, M., Goh, L., Vallegra, F., Bouffanais, R., & Yue, D. K. P. (2018). Distributed system of autonomous buoys for scalable deployment and monitoring of large waterbodies. *Autonomous Robots*, 42(8), 1669–1689. <https://doi.org/10.1007/s10514-018-9702-0>