

Balancing exploitation of renewable resources by a robot swarm

Roman Miletitch · Marco Dorigo · Vito Trianni

Received: date / Accepted: date

Abstract Renewable resources like fish stock or forests should be exploited at a rate that supports regeneration and sustainability—a complex problem that requires adaptive approaches to maintain a sufficiently high exploitation while avoiding depletion. In presence of oblivious agents that cannot keep track of all available resources—a frequent condition in swarm robotics—ensuring that the exploitation effort is correctly balanced is particularly challenging. Additionally, the possibility to exploit resources by multiple robots opens the way to focusing the effort either on a single or on multiple resources in parallel. This means that the swarm needs to collectively decide whether to remain cohesive or split among multiple resources, as a function of the ability of the available resources to replenish after exploitation. In this paper, we propose a decentralised strategy for a swarm of robots that adapts to the available resources and balances the effort among them, hence allowing to maximise the exploitation rate while avoiding to completely deplete the resources. A detailed analysis of the strategy parameters provides insights on the working principles and expected performance of the robot swarm.

Keywords swarm robotics · resource exploitation · foraging · load balancing

1 Introduction

Exploration of the environment and exploitation of valuable resources represent problems commonly studied in robotics, and in particular in the multi-robot systems and swarm robotics sub-fields (Winfield, 2009; Ducatelle et al.,

R. Miletitch · M. Dorigo
IRIDIA, Université Libre de Bruxelles, Brussels, Belgium
E-mail: rmiletit@ulb.ac.be, mdorigo@ulb.ac.be

V. Trianni
ISTC - Italian National Research Council, Rome, Italy
E-mail: vito.trianni@istc.cnr.it

2014; Trianni and Campo, 2015). The abilities of a robot to move in space and to autonomously identify locations of interest make this problem particularly relevant for a number of different application scenarios, from search and rescue to mining, from precision agriculture to space exploration (Murphy et al., 2008; Cheein and Carelli, 2013; Yoshida, 2009; Trianni and Dorigo, 2005). Exploration and exploitation are activities commonly executed in the animal kingdom as well, be it to search for mating partners, prey, or other valuable goods. Differently from artificial situations, natural examples often lead to some ecological equilibrium in which exploitation remains constrained within the environment *carrying capacity* (Hui, 2006). This means that natural systems evolved adaptive strategies to balance the exploitation effort among the available resources (e.g., through dietary diversification, see Simpson et al., 2004), so that resource depletion is avoided and survival is granted for the whole population. Indeed, when resources are renewable, greedy strategies do not always pay off as the resource depletion would mean losing an important income and expending energy for further exploration. For this reason, animals exploit strategies like trap-lining, which allows to iterate among known resources hence granting sufficient time for replenishment (Saleh and Chittka, 2006). If, however, cognitive limitations or other ecological constraints do not support the development and utilisation of long-term memory about resource locations, stochastic search strategies remain the only viable alternative (Bartumeus et al., 2005).

In this study, we propose an adaptive strategy to balance exploitation of renewable resources by a robot swarm. Individual limitations, such as the inability of robots to be aware of each available resource and its profitability, as well as the fact that their behaviour may be constrained to following simple reactive rules, entail that achieving a correct balance is not trivial. In such conditions, a balanced exploitation should result from a collective self-organising process in which information about the availability of resources is shared among the robots to achieve a correct allocation, preventing resource depletion and maximising the flow of goods. To this end, we draw inspiration from the behaviour of honeybees, and exploit the guidelines provided by a design pattern hinged on their nest-site selection behaviour (Reina et al., 2015b,a). Despite being conceived for collective decision making, the design pattern indicates how the swarm dynamics can switch between convergence toward the exploitation of a single resource (when its quality is good enough to sustain a large swarm) and balancing between the available resources (when no resource has sufficient quality to sustain the whole swarm).

In this paper, we focus on a resource exploitation problem whereby items have to be collected from multiple resources, which can replenish at a fixed, unknown rate. We borrow and adapt the collective decision-making algorithm from Reina et al. (2015b), and we improve on it by studying for the first time how this algorithm adapts to dynamic environmental conditions that result from continuous resource replenishment after foraging. Indeed, the variation of the number of items within a single resource corresponds to a variation of its perceived quality, which has an impact on the macroscopic dynamics of the

proposed algorithm that has never been studied to date. To this end, we isolate the different processes determining the collective dynamics and study their impact. This contributes to the identification of a parameterisation that can lead to a balanced exploitation not only with respect to the regeneration rate of the resources—hence avoiding resource depletion—but also with respect to their distance from the central place where foraged items need to be retrieved. Integrating all these aspects proves particularly challenging, and we present here a large-scale study on the most important parameters determining the system behaviour.

The paper is organised as follows. In Section 2, we discuss the state of the art in swarm robotics for exploration and resource exploitation problems. In Section 3, we describe the experimental setup. The results obtained in different experimental conditions are presented in Section 4, where the parameters governing the swarm behaviour are analysed to identify the most suitable values to optimise the resource collection rate and exploitation efficiency. Discussions and conclusions are drawn in Section 5.

2 Background

The exploration and resource exploitation problem has been previously approached in swarm robotics, mainly for non-renewable resources. Several studies present adaptive foraging algorithms inspired by the well-known response thresholds model (Bonabeau et al., 1996). Krieger et al. (2000) studied the effects of heterogeneities in the individual response thresholds and of additional recruitment mechanisms to adapt the size of the foraging group to the features of the available resources. Labella et al. (2006) tested an extended response threshold model with individual learning abilities in a group of robots that foraged for sparse resources, and observed the adaptation of the robot activities between foraging and idleness, linking it to hardware differences among robots. Liu et al. (2007) employed a similar adaptation mechanism to allocate workers for a foraging task, and later presented a macroscopic probabilistic model that predicts the robotic system dynamics (Liu and Winfield, 2010). An adaptive response threshold model was presented by Castello et al. (2015), tailored to fast adaptations to changing environmental conditions. In the above mentioned studies, task allocation resulted from the adaptivity of the individual behaviour, which balances the foraging rates on the basis of information collected about the resource availability. However, resources were sparsely available in the environment, hence limiting the need for information sharing about the resource locations.

When resources are clustered in specific areas in space, recruitment of robots to areas in which resources are likely to be found becomes important (Krieger et al., 2000). Gutiérrez et al. (2010) studied the collective behaviour of robots foraging from static resources and sharing information about the resources position, eventually leading to the exploitation of the closest one thanks to a positive feedback given by a larger number of robots promoting

the closer alternative. A similar study was presented by Miletitch et al. (2013), in which the swarm performance was affected by the way in which information received from peers was integrated with the individual knowledge. Hecker and Moses (2015) developed a foraging algorithm based on a delicate balance between individual search and recruitment from peers, and optimised the system parameters through a genetic algorithm to fit different environmental conditions, including clustered resources. Similarly, Pitonakova et al. (2016) considered foraging of resources possibly clustered in various deposits, also taking into account dynamic conditions where the quality of the deposit abruptly changed, to evaluate the plasticity of the proposed behaviour. In similar conditions, the need to select the most profitable resource among many available can lead to collective decision-making problems (Valentini et al., 2017).

The studies mentioned so far did not deal with sustainable resource exploitation, but instead optimised the foraging efficiency, either by choosing the most profitable resource or by switching to different resources when the current one gets depleted. Sustainable or continuous foraging was instead focused on the optimisation of the foraging rate, and in the maintenance of resources in lieu of their depletion (Song and Vaughan, 2013; Liemhetcharat et al., 2015). The “maximum sustainable yield model” introduced by Song and Vaughan (2013) prescribes that resources characterised by a logistic growth should be maintained at the level of maximum regeneration rate. An algorithm was proposed to allocate a slightly higher number of robots to each resource, where each robot adapted its foraging rate to maintain the resource around the optimal size for regeneration. Maximisation of the foraging rate was also studied by Liemhetcharat et al. (2015), who however employed an heterogeneous system in which some agents could have an overview of the resource exploitation level, and helped the other foraging agents to adjust their activity rate so as to maximise the system efficiency. Improving on these studies, we present an algorithm for resource exploitation by a swarm of homogeneous robots that does not require global information, and that can tune the exploitation effort on the basis of the ability of individual robots to successfully forage from a given resource. Differently from what happens in adaptive response threshold approaches (Labella et al., 2006; Castello et al., 2015), here the exploitation effort is adjusted without any individual learning component. Similarly to the study by Pitonakova et al. (2016), our study seeks a balance between exploration and exploitation of resources, [but here we consider a continuous change of resource attractiveness that results from the foraging dynamics and from the spontaneous resource replenishment.](#)

3 Exploration and exploitation of resources by robot swarms

3.1 Problem description

This study focuses on a sustainable resource exploitation problem whereby robots have to search for items scattered in an open environment and re-

trieve them to a home location (hereafter referred to as “nest”, in analogy to foraging by social insects). The nest is a circular area (radius: 0.8 m) at the center of the robot arena represented by a black disk painted over an otherwise white floor (see Figure 1). Retrievable items are cylinder-shaped objects (radius: 0.05 m) clustered together to form a “resource”. In this study, we focus on a simplified exploitation problem in which only two resources are present—labelled A and B —although the proposed solution can be easily generalised to larger numbers of resources. Each resource $i \in \{A, B\}$ contains at most $M_i = 30$ items positioned according to a 2D Gaussian distribution around the resource centre ($\sigma_R = 0.35$ m), keeping a minimum distance $d_{min} = 0.14$ m between items. Resources are characterised by quality and position. The resource quality r_i is defined by the rate of creation of new items ($r_i \in \{0.01, 0.03, 0.05, 0.1\}$ items/second). The resource position is defined by the distance from the nest ($d_i \in \{4, 6, 8, 10\}$ m). The relative angle α_{AB} between resources is chosen at random with a minimum angle of $\frac{\pi}{3}$ between different resources to ensure separation. When the number of items within resource i is lower than M_i , new items are generated with the given regeneration rate r_i , practically implementing a Bernoulli model for resource regeneration as described by Liemhetcharat et al. (2015).

3.2 Robots and simulations

Experiments are run in simulation using ARGoS (Pinciroli et al., 2012). The simulator models the *marXbot* robot, a differential drive robot developed for swarm robotics research (Bonani et al., 2010; Dorigo et al., 2013). MarXbots have a differential drive motion, and their speed is measured by a wheel encoder simulated to have a gaussian noise of 5%, which supports simple odometry and local positioning/navigation. Avoidance of both items and other robots is done at short range (≈ 10 cm) with infra-red proximity sensors. Robots can avoid each other at a longer range (≈ 1 m) exploiting the infrared range and bearing system (Roberts et al., 2009), which is also used for communication among robots. [Information from long range sensors is exploited to compute a repulsive virtual force that pushes the robot away from neighbours \(Borenstein and Koren, 1989\):](#)

$$\mathbf{V} = \sum_{i \in \mathcal{N}} \frac{D_M - |\mathbf{v}_i|}{D_M} e^{-i \angle \mathbf{v}_i} \quad (1)$$

where \mathcal{N} is the set of robots within the maximum considered distance D_M , and \mathbf{v}_i is a vector indicating the relative position of robot $i \in \mathcal{N}$. Each robot contributes to the repulsive force \mathbf{V} with a component that is inversely proportional to distance, and in the opposite direction from the robot. The obstacle avoidance behaviour has been optimised to minimise the effects of robot density and congestion on the ability to navigate back and forth between resources, as detailed in a previous study (Miletitch et al., 2013). Robots recognise the nest by means of infra-red ground sensors, which allow robots to differentiate

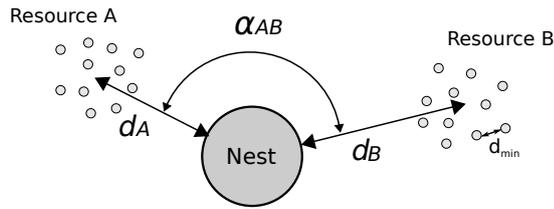


Fig. 1 The resource exploitation problem requires robots to bring cylindrical items back to the nest. The experimental setup defines two resources—labelled as resource A and B —positioned at distance d_A and d_B from the nest, and separated by the angle $\alpha_{AB} > \frac{\pi}{3}$. Each resource i contains at most M_i items, scattered around the resource center according to a 2D Gaussian distribution, keeping a minimum distance d_{min} between items. The nest is a circular area (radius: 0.8 m) painted in black to be recognised through the robots infrared ground sensors.

between the white color of the floor and the black color of the nest. Furthermore, robots can localize the items to be collected using the omnidirectional camera to perform a simple blob detection, and are able to recognise items up to a distance of 1 m. Items can be reached and then grabbed with a specific claw that rotates around the robot, making it easier to navigate while holding an item. Robots can communicate with each other by means of the infrared range and bearing system (Roberts et al., 2009), which allows to locally broadcast short messages that can be perceived within a distance of approximately 0.7 m. Robots broadcast at regular intervals of 0.1 s, and there is no re-broadcast of information received (no multi-hop communication).

3.3 Individual and Collective Behaviour

The overall goal of the robot swarm is to maximise the retrieval rate, that is, the number of items per unit time that are successfully transported to the nest. We assume that robots have no a priori knowledge about the position and profitability of resources, neither they have any map of the environment to support navigation. Hence, exploration is required to gather information on the available resources. We also assume that robots can track—e.g., through odometry—the position of the nest and of at most one resource at the time, which is the one most recently visited. In this way, we ensure that robots have up-to-date information about the state of the resources they have found, avoiding to memorize the location of resources that may be unprofitable or completely depleted. Finally, we assume that robots can carry only one item at a time, hence multiple robots can forage from the same resource at the same time to maximise the exploitation rate.

The desired swarm behaviour requires a recruitment process, so that robots can spend less time in exploration, focus on the available resources and exploit them in parallel. When multiple resources are available, exploitation can be focused on one of them if its quality is high enough to sustain the whole swarm. Otherwise, a balanced exploitation of the two available resources is

preferred. This collective behaviour has properties similar to value-sensitive decision-making studied in house-hunting honeybees (Pais et al., 2013; Reina et al., 2017). Indeed, when engaged in a collective decision, a honeybee swarm may arrive at consensus when the quality of the option is sufficiently high, otherwise remain in a “undecided” state when the quality is low, in the hope that a better alternative is discovered later. As a matter of fact, such undecided state corresponds to the swarm being split in sub-populations committed to the low-quality alternatives that they could find. This “undecided” state can be seen as a load-balancing state, because the house-hunting swarm is split among potential nest-sites, much as the foraging swarm is split among the available resources. This highlights the usefulness of framing the load balancing problem studied here in terms of a value-sensitive collective decision problem.

Starting from this observation, we decided to synthesise the individual robot behaviour taking inspiration from a design pattern derived from the honeybee nest-site selection behaviour (Reina et al., 2015b). The design pattern provides guidelines to implement the individual behaviour as a probabilistic finite state machine (PFSM), where any robot can be in two macro states: committed to exploit a known resource, or uncommitted and exploring. Additionally, upon robot-robot encounters, local information exchange can lead to changes in the commitment state. Overall, four concurrent processes need to be implemented in the individual behaviour: (i) spontaneous discovery of any resource i with probability $P_{D,i}$; (ii) spontaneous abandonment of commitment to resource i with probability $P_{L,i}$; (iii) recruitment of uncommitted agents following interaction with a robot committed to resource i with probability $P_{R,i}$; (iv) inhibition of commitment, whereby an agent committed to resource i becomes uncommitted after interaction with a robot committed to resource $j \neq i$, with probability $P_{I,j}$ (cross-inhibition). These probabilities are either completely defined by the problem itself (e.g., $P_{D,i}$ for discovery of resource i , see Section 4.1) or are parameters defined at design time in order to tune the collective behaviour and to achieve the desired exploitation of the available resources. Cross-inhibition is particularly relevant, as it can determine the switch from the parallel exploitation of multiple resources to full convergence towards a single resource. This mechanism has been observed in house-hunting honeybees (Seeley et al., 2012), and is used to adaptively select nest sites of high quality, quickly abandoning those of low value (Reina et al., 2017). Indeed, through cross-inhibition, agents committed to a resource can return uncommitted, explore for other—possibly better—alternatives, or get recruited by other agents.

In this work, we have implemented the individual behaviour as the PFSM represented in Fig. 2, which is executed every $\Delta t = 0.1$ s. Here, boxes represent macro states corresponding to the commitment state of a robot, while circles represent micro states corresponding to basic behaviours executed until some (probabilistic) transition is triggered. The robot is considered to be committed to a resource when it knows its location, otherwise it is considered uncommitted. The actual movements of the robot are governed by the following basic behaviours:

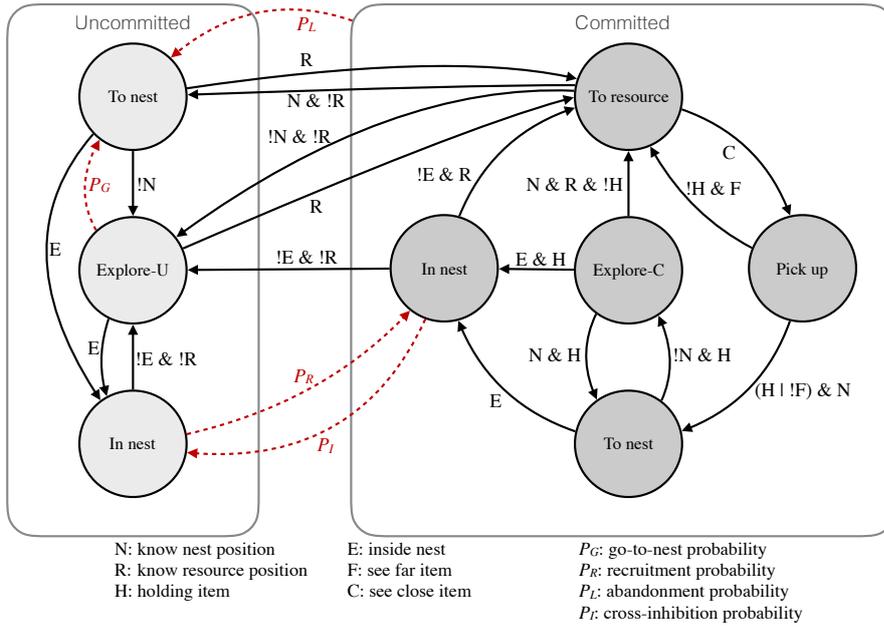


Fig. 2 Probabilistic finite state machine (PFSM) representing the individual robot behaviour. The two boxes represent macro states for the committed and uncommitted robot. Circles inside the macro states represent PFSM states in which a robot executes a basic behaviour. The “Explore” behaviour is executed in different conditions depending on the commitment state, hence states are named differently to avoid confusion. Arrows represent transitions between states, and are triggered when a certain Boolean expression is verified (AND: ‘&’; OR: ‘|’; NOT: ‘!’; see legend at the bottom). When a robot is uncommitted (left box), it has no knowledge of any resource and searches for it, periodically returning to the nest. When committed (right box), the robot knows where a resource is and tries to retrieve items from it. The red dashed arrows represent probabilistic transitions.

- **Explore**: in this state, the robot explores the arena performing a correlated random walk (Dimidov et al., 2016). Whenever sufficient information becomes available (e.g., location of nest and resources), either through exploration or following interactions with other robots, a different behaviour may be triggered.
- **To resource**: the robot moves toward the location of a known resource to search for more items to retrieve.
- **Pick up**: when some item is in close range, the robot navigates toward it and picks it up. Should the grasping procedure fail, the robot tries again or chooses another item to pick up, if available.
- **To nest**: the robot navigates back to the nest, possibly bringing back an item to deposit.
- **In nest**: when in the nest, the robot deposits the item it is carrying—if any—and then performs a random walk until it moves out of the nest.

Robots start from the nest at the beginning of each experiment and keep track of their positions through individual and social odometry (see Gutiérrez et al., 2010, more details below).

When a robot is uncommitted, it explores the environment to gather information about the location of the nest (if unavailable) and of the resources (see left box in Fig. 2). With a fixed probability P_G , a robot stops exploring and returns to the nest, where it has a high probability of interacting with other robots. With this mechanism, exploration is constrained around the nest and robots do not wander away for too long.

When a robot is committed, it moves back and forth between the nest and a known resource to retrieve some item (see right box in Fig. 2). If it loses track of the nest, it explores the neighbourhood until either it finds the nest or it receives its location by other robots in its neighbourhood. When in the nest, it deposits the carried item and starts another exploitation trip to the known resource. At any time—and from any state in the committed macro state—a robot can abandon commitment for resource i with probability $P_{L,i}$. This corresponds to erasing the information about the resource and returning back to the nest, from where to retrieve exploration.

Robots interact locally through infrared communication, broadcasting at regular intervals their knowledge about the position of the nest and of the known resource, if available. Such information is used for two purposes. On the one hand, it is used by neighbours to update their own knowledge about the same locations, following the social odometry paradigm (see Gutiérrez et al., 2010, for details). This assures that the swarm maintains through time a good overall knowledge of the nest and resource positions (Miletitch et al., 2013). When robots are located in the nest, the same message can lead to recruitment and cross-inhibition of uncommitted and committed robots, respectively. The uncommitted robot can get recruited with a probability $P_{R,i}$ by another robot committed to resource i upon reception of a message. Similarly, a robot committed to resource i can get inhibited with probability $P_{I,j}$ —and turns uncommitted—upon reception of a message from a robot committed to resource $j \neq i$ (see Fig. 2).

In this study, the probability of discovering a resource results from the random exploration that robots perform, and is dependent on the distance d_i of resource i from the nest: the closer the resource, the higher the probability of discovering it (see Section 4.1). On the other hand, the other probabilities introduced are control parameters that determine the overall macroscopic behaviour of the robots. Here, we use fixed probabilities independent of the resource quality, hence $P_{L,i} = P_L$, $P_{R,i} = P_R$ and $P_{I,i} = P_I$, and we perform a thorough analysis to uncover the effects of the control parameters on the emergent swarm behaviour.

4 Results

As mentioned above, our goal is to study the macroscopic behaviour resulting from the rules defined in Sect. 3.3 for different values of the control parameters we identified. We want to obtain different types of macroscopic behaviour, from exploitation of a single good resource to a load balancing between two different resources. Additionally, we want to maximise the exploitation efficiency of the robot swarm by optimising the rate of retrieved items, either from one or from multiple resources. To understand the effects of the different processes determining the collective dynamics, we performed a set of experiments to isolate the contribution of each component of the developed behaviour. In Sect. 4.1, we analyse the exploration efficiency when robots are uncommitted, while in Sect. 4.2 we focus on the exploitation efficiency when robots are committed to a given resource. We analyse the effects of recruitment in determining the tradeoff between exploration and exploitation in Sect. 4.3. Then, we report on the effects of cross-inhibition on the ability of a swarm to converge to the exploitation of a single resource or split among multiple ones, and we assess the exploitation efficiency of the swarm when dealing with multiple resources (see Sect. 4.4).

4.1 Baseline exploration efficiency

Resource exploration is the activity that agents perform when uncommitted. As mentioned above, robots perform a correlated random walk (Dimidov et al., 2016) until they find an item to be picked up. Random walk continues either until a resource is found, or until the robot is triggered (with probability P_G) to return home. When in the nest, a robot can share information about the resource found, or interact with other robots.

To evaluate the exploration efficiency of the swarm, we run a series of experiments to measure (i) the average rate of discovery of a resource with respect to the distance, and (ii) the average percentage of robots that are found in the nest. The former gives us an idea of the probability of discovery P_D as a function of the distance of a resource from the nest: the higher this probability, the sooner the swarm can start exploiting a resource. The latter gives us an idea of the ability of robots to interact with each other when uncommitted, and has a bearing on the ability to be recruited by other robots to a known resource. Both metrics depend on the probability P_G and on the distance d_i of resource i from the nest.

Experimental Setup In this set of experiments, we use $N = 40$ robots that are constrained to remain in the uncommitted state: whenever a resource is found, a robot goes back to the nest, but does not store the resource location. In this way, when in the nest a robot starts again an exploration trip. We provide only one resource to be found with high regeneration rate $r = 0.1$ items/s, so that it remains close to the maximum number of items (i.e., 30 items).

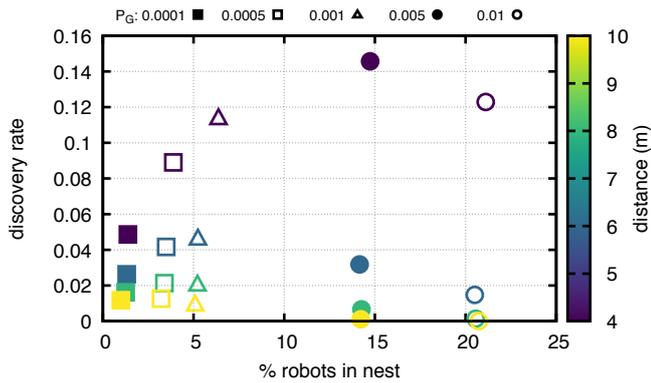


Fig. 3 Exploration efficiency: We highlight the tradeoff between the average rate of discovery of a resource and the average percentage of robots found within the nest. Lighter points correspond to larger distances, varying between 4 m and 10 m. Different point types correspond to different values of P_G .

This resource is placed at a fixed distance $d \in \{4, 6, 8, 10\}$ m from the nest. Additionally, we vary the probability to spontaneously go back to the nest $P_G \in \{0.0001, 0.0005, 0.001, 0.005, 0.01\}$. These values have been chosen to ensure a sufficient searching time before returning to the nest. Considering that the probabilistic choice is taken 10 times per second, the average searching time corresponds to $(10P_G)^{-1}$ s. We simulate the exploration for $T = 2000$ s, and we measure the rate of discovery and the percentage of robots found in the nest, in average.

Results Figure 3 summarises the average results from 100 runs performed in each condition, varying resource distance and probability to go back to the nest. It is possible to note an expected pattern for which, the higher the distance of the resource, the smaller is the discovery rate (see the color shades of the different points). Similarly, the lower the probability P_G , the smaller the percentage of robots in the nest. The distance of the resource also has an impact, although relatively small, on the percentage of robots found in the nest, because robots take less time to travel from the resource to the nest once the resource is found. Indeed, such a shift in the percentage of robots within the nest is visible especially for higher discovery rates.

Overall, we note that resources that are 10 m away from the nest are difficult to discover, and only small enough values for P_G ensure a non-null rate. However, P_G should not be too small, in order to grant a sufficient percentage of robots within the nest. A suitable tradeoff is given by $P_G = 0.001$, and we choose this value for the following experiments. With this value, the rate of discovery is non-null also for large distances, and at the same time the percentage of robots within the nest remains reasonably high, allowing for sufficient robot-robot interactions.

4.2 Baseline efficiency in resource exploitation

Once the exploration efficiency of uncommitted robots has been determined, we can evaluate the exploitation efficiency of committed robots. This way, we study separately both sub-behaviors (committed and uncommitted) of the implementation. Generally speaking, we can define the exploitation efficiency of a swarm as the overall rate of retrieval of items, that is, the number of items retrieved by all robots per second independently of the resource from which the items were collected. The retrieval rate will depend on the quality of the resource, the distance from the nest, and the number of robots committed to the resource and actively collecting items from it. Clearly, the way in which the individual behaviour is implemented could lead to interferences and congestion that have an impact on the overall retrieval rate. In order to evaluate the maximum efficiency of a swarm given the implemented behaviour (i.e., navigate back and forth from resources and pick up and deposit collected items, see Sect. 3.3), we perform a series of experiments largely varying the experimental conditions. We then introduce a model of exploitation of multiple resources which provides a baseline to evaluate the overall efficiency of the swarm when decision making and load balancing will be introduced.

Experimental Setup We consider the case in which a fixed number $N \in [1, 40]$ of robots exploit a single non-depletable resource (i.e., a resource with the maximum regeneration rate $r = 0.1$ items/s containing at most 30 items), placed at a fixed distance from the nest ($d \in \{6, 8\}$ m). We measure the retrieval rate of a group of robots that continuously exploit the resource. To this end, we force robots to stay committed to the given resource and we provide them with perfect information about the resource location (i.e., robots never lose track of the resource and can always navigate back and forth between resource and nest). Under these conditions, we measure the overall rate of returned items per second once its evolution reaches a plateau.

Results Figure 4 summarises the average results obtained from 100 independent runs in each experimental condition, varying N and d_i . Each run lasts $T = 2000$ s during which robots continuously exploit the known resource. In the left panel of Fig. 4, the retrieval rate $R(d, N)$ is shown, indicating a linear dependency between group size N and rate of retrieval. This implies that, for the group sizes and distances considered, there is no negative impact from interferences or congestion, which would instead result in a sub-linear growth. We also note, as expected, that the efficiency is higher for closer resources, due to the fact that robots need to travel shorter distances. This will have an impact on the collective behaviour when multiple resources are presented at the same time. Indeed, the swarm may have to face the choice between exploiting a close-but-poor resource, or a farther-but-rich one. The correct balance between the two must emerge from the different expected efficiencies in the exploitation.

To evaluate the efficiency in presence of two resources with different quality and distance, we provide a compact visualisation built on top of the maximum efficiency computed for single resources and fixed groups. We consider here a total group size of $N = 40$ robots, and we compute the expected efficiency for all possible allocations of robots to resources A and B , under the assumption that $N = N_A + N_B + N_U$, where N_U is the number of uncommitted robots, which therefore do not contribute to the exploitation. Given the maximum retrieval rate $R(d, N)$, experimentally obtained for non-depletable resources at different distance shown in Fig. 4 left, we compute the normalised exploitation efficiency of a resource of quality r as follows:

$$\hat{R}(d, N, r) = 1 - \left| \frac{R(d, N) - r}{R(d, N) + r} \right| \quad (2)$$

which has its maximum when $R(d, N) = r$, corresponding to a resource that can completely support exploitation from N robots without being depleted. The normalised exploitation efficiency slightly decreases when $R(d, N) > r$, corresponding to the over-exploitation of the resource, leading to complete depletion. We use this simple model to visualise the combined efficiency in foraging from two resources in parallel. We show the combined efficiency as an heatmap on a ternary plot (see the right panel of Fig. 4). Here, each point $\langle N_U, N_A, N_B \rangle$ corresponds to a given allocation of robots to the two available resources A and B . We color-code the normalised retrieval rate \hat{R}_A in shades of red (see the horizontal axis for $N_B = 0$), while \hat{R}_B is visualised in shades of green (see the vertical axis for $N_A = 0$). The combined efficiency is rendered as the sum of the two colours, hence bright yellow for the optimal values given the resources' quality and distance (see Fig. 4 right for an example). This visualisation allows to indicate if a certain allocation of robots corresponds to

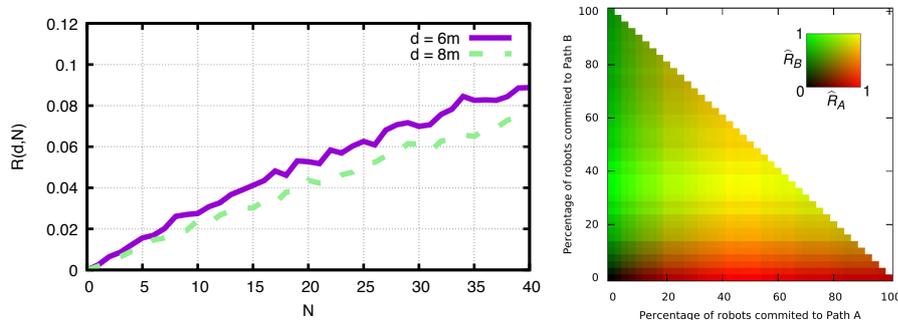


Fig. 4 Left: exploitation rate $R(d, n)$, computed as the number of items per second retrieved from a non-depletable resource at varying distance d and with varying group size N . Right: Combined efficiency in the exploitation of two resources with different quality and distance. In this example, we consider the case for $d_A = 6, d_B = 8, r_A = 0.03, r_B = 0.05$. The normalised exploitation efficiency \hat{R} is represented in red for resource A and in green for resource B . The combination of both—as indicated by the inset—represents the total efficiency: dark zones imply low overall efficiency and bright ones imply high efficiency. Yellow zones (as a mix of green and red) imply that both zones are exploited in parallel.

the balanced exploitation of both resources, and will be used to evaluate the actual efficiency of the swarm when two resources are present, as discussed in Sect. 4.4.

4.3 Exploration vs. exploitation of a single resource

The tradeoff between exploiting a given resource and exploring in search of other possibilities is the result of a delicate balance of multiple forces, and needs to be carefully studied. With the implemented behaviour, robots commit to a resource either upon discovery through random search, or upon recruitment from an already committed robot. While individual discoveries are always occurring with a constant probability, as shown in Sect. 4.1, the probability of a robot to be recruited grows with the size of the recruiting population. This creates a positive feedback loop for which, the more a swarm exploits a specific resource, the more it recruits to it. On the other hand, resource depletion following excessive foraging provides a negative feedback that works in the opposite direction and tends to stabilise the system, because those robots that do not find an item when they reach a depleted resource turn uncommitted and stop recruiting once back to the nest. To evaluate the coupled effects from recruitment and over-exploitation, we analyse the dynamics of a swarm presented with a single resource of varying quality.

Experimental Setup We consider the case of a single resource of quality $r \in \{0.01, 0.03, 0.05, 0.1\}$ items/s, placed at distance $d = 8$ m from the nest. Here, the number of items within a resource can decrease so that resources can get depleted upon high exploitation. Robots execute the complete behaviour discussed in Sect. 3.3, although cross-inhibition is not present as there is only a single resource. We consider a constant probability to go back to the nest $P_G = 0.001$ as resulting from the experiments discussed above. Coherently, we fix the probability of abandonment to $P_L = 0.001$ to have a similar rate of abandonment in both exploration and exploitation. When in the nest, robots can only recruit each other with a probability $P_R \in \{0.01, 0.02, 0.03\}$. Also in this case, every run is executed for $T = 2000$ s using $N = 40$ robots, and we perform 100 experimental runs for each experimental condition.

Results To appreciate the macroscopic dynamics resulting from different parameterisations, we show the average percentage of robots committed to the resource, the fraction of robots that switch commitment state per second, and the fraction of items left in the resource (see Figure 5). When the quality is high ($r = 0.1$), nearly all robots get committed to the resource and exploit it, and the number of items remaining in the resource stays very high ($> 80\%$). Small recruitment probabilities correspond to a slow increase of the committed population, until a plateau is reached. The larger values we tested lead to a much quicker increase of the committed population, which stabilises earlier (see the top row in Fig. 5). Lower qualities of the resource ($r \leq 0.05$) lead to a

balance between positive and negative feedbacks that stabilises the committed population to a value that strongly depends on the resource quality r , and to a much lesser extent also on the recruitment probability P_R . Interestingly, we also observe a higher rate of change in the commitment state in correspondence of higher values of P_R , which suggests that the macroscopic dynamics oscillate around the average values displayed in Fig. 5. Overall, the main impact of P_R is on the speed of growth of the committed population. Fast growth is useful for high-quality resources, but not so much for low quality ones, as the risk to quickly over-exploit the resource may lead to fast depletion of the resource and strong instabilities and oscillations due to massive abandonment. Hence, we consider a value of $P_R = 0.02$ as suitable for balancing quick growth with stability of exploitation.

4.4 Balancing resource exploitation

Whenever two or more resources are present, robots need to choose which resource to exploit, and a competition between the sub-populations committed to the one or the other resource is observable due to committed robots

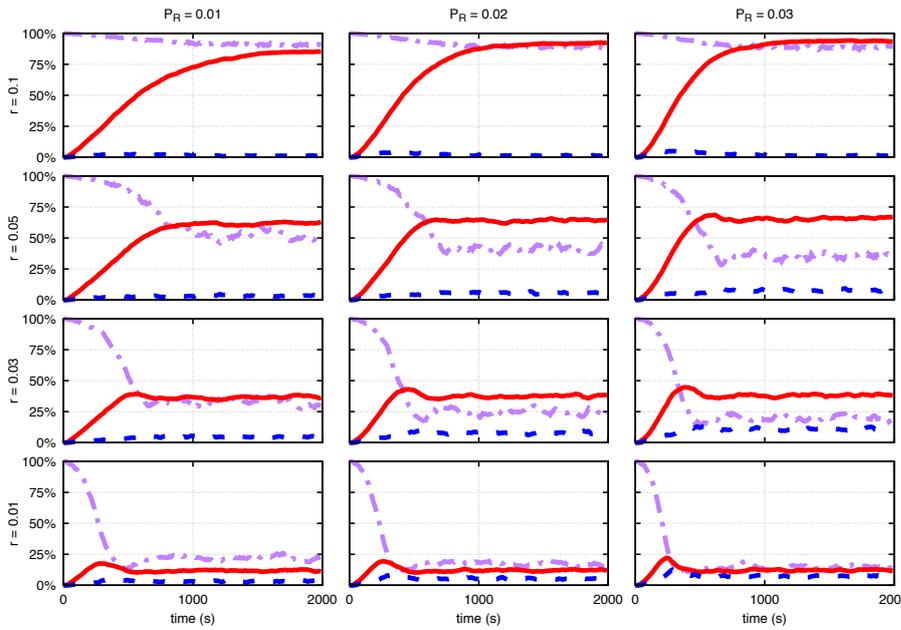


Fig. 5 Dynamics of exploitation vs exploration with varying resource quality and recruitment probability. Each graph shows the variation through time of the percentage of (i) committed robots (solid red lines), (ii) robots switching commitment state (dashed blue lines) and (iii) items left in the resource (dot-dashed purple lines). Each graph represents a different setup with a distance of $d = 8$ m, recruitment probability $P_R \in \{0.01, 0.02, 0.03\}$ and resource quality $r \in \{0.01, 0.03, 0.05, 0.1\}$

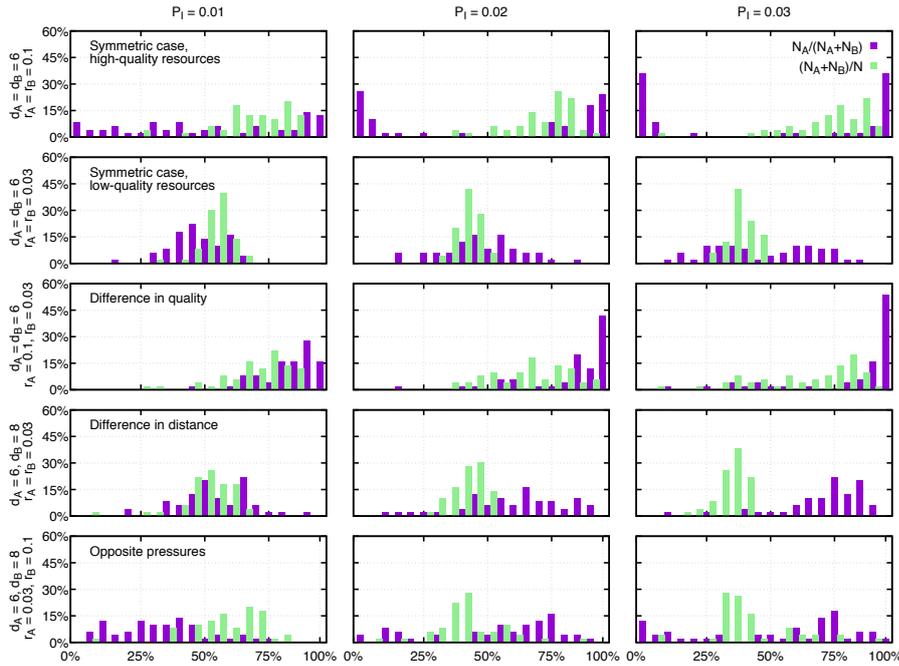


Fig. 6 Allocation of robots to different resources. Each histogram represents the distribution of observables across $N = 100$ independent runs. Light green bars represent the percentage of committed robots $(N_A + N_B)/N$. Dark violet bars represent the percentage of committed robots that have chosen resource A: $N_A/(N_A + N_B)$. Each plot represents a different setup, defined by the value of P_I and the quality and distance of the two resources. The experimental setup characterising each row is detailed on the left. The value of P_I is reported on top of each column. For additional results in different experimental conditions, see the figures in the supplementary material.

recruiting uncommitted ones and cross-inhibiting each other. To evaluate the extent to which such competition leads to a balanced exploitation, we run a set of experiments to understand what is the average allocation of robots among committed populations as a function of varying resources' quality and distance, and for different values of the cross-inhibition probability P_I .

Experimental Setup In this experiment, we consider two available resources which can be at varying distance $d_i \in \{6, 8\}$ m and varying quality $r \in \{0.01, 0.03, 0.05, 0.1\}$ items/s. Robots execute the complete behaviour presented in Sect. 3.3, with $P_G = P_L = 0.001$ and $P_R = 0.02$, in accordance to the experiments presented above. Here, robots committed to different resources can cross-inhibit each other with probability $P_I \in \{0.01, 0.02, 0.03\}$. We perform 100 runs that last each $T = 2000$ s, and we look at the final allocation of robots committed to the different resources, or uncommitted. Additionally, we discuss the overall efficiency of the exploitation of the two resources in parallel, following the empirical model introduced in Sect. 4.2.

Results Overall, cross-inhibition defines how tolerant the swarm is of having a segmented population: the smaller the cross-inhibition probability, the lower the negative interaction between committed populations, the higher the probability that sub-populations committed to different resource can coexist (see also similar results from the macroscopic models by Pais et al., 2013; Reina et al., 2017). The resulting dynamics can lead to convergence to a single resource or balancing among many. This can be understood by looking at the distribution of the commitment state of the robots at the end of each run. The histogram in Fig. 6 represent such distributions, specifically for the percentage of robots committed to any resource (that is, $\frac{N_A+N_B}{N}$) and for the percentage of committed robots that chose resource A (that is, $\frac{N_A}{N_A+N_B}$). The former informs us about the ability of robots to successfully exploit resources, in average, given the experimental conditions. The latter informs us about the tendency of committed robots to choose resource A (and conversely to not choose resource B), hence revealing the collective choice or load balancing achieved by the swarm. In our experiments, we observe the full range of macroscopic dynamics for varying experimental conditions, which we discuss in the following (see Fig. 6).

Symmetric case, high-quality resources ($d_A = d_B = 6$ m, $r_A = r_B = 0.1$ items/s, first row in Fig. 6). In this condition, both resources could support the whole swarm, hence a collective decision in which the swarm achieves convergence on the exploitation of a single resource may lead to the best results, as in this condition there are practically no robots uncommitted (i.e., light green bars are shifted to high percentages). We observe that low values of P_I do not grant convergence, resulting in a uniform repartition of robots over the two resources in the different runs (see the dark violet histograms). Higher values of $P_I \geq 0.02$ result instead in a collective decision, as observable from the bi-modal distribution of $N_A/(N_A + N_B)$, indicating full commitment for either A or B .

Symmetric case, low-quality resources ($d_A = d_B = 6$ m, $r_A = r_B = 0.03$ items/s, second row in Fig. 6). In this condition, no resource can sustain the whole population, and we therefore observe a somewhat equal allocation of the committed robots among the two resources, especially for low cross-inhibition values ($P_I \leq 0.02$) where a unimodal distribution is present (see dark violet bars). Stronger cross-inhibition leads to a large competition and the appearance of a bi-modal distribution, although less pronounced than in the previous case. The number of uncommitted robots is in general high due to the low quality of the resources (corresponding to light green bars centred around low percentages of committed robots), and increases for larger values of P_I .

Difference in quality ($d_A = d_B = 6$ m, $r_A = 0.1$ items/s, $r_B = 0.03$ items/s, third row in Fig. 6). In this condition, resources differ only in the rate of replenishment, leading to a bias towards the choice of the most profitable one (A in this case, see the dark violet bars shifted towards high percent-

ages). The higher the cross-inhibition probability, the stronger the shift of the distribution toward the high-quality resource.

Difference in distance ($d_A = 6$ m, $d_B = 8$ m, $r_A = r_B = 0.03$ items/s, fourth row in Fig. 6). In this case, resources are both of somewhat low quality, but one is farther away than the other, leading to an exploitation balancing biased towards the closer resource. We can observe here the high number of uncommitted robots, due to the low quality and the large distance of one resource (see the distribution of committed robots centered around low percentages). The effects of the cross-inhibition probability are largely similar to the symmetric case with low-quality resources (second row), but the distribution is biased toward the closer resource (A) especially for larger values of P_I .

Opposite pressures ($d_A = 6$ m, $d_B = 8$ m, $r_A = 0.03$ items/s, $r_B = 0.1$ items/s, last row in Fig. 6). This condition represents the most difficult case for the algorithm as the asymmetries in distance and quality oppose and may compensate each other. Indeed, the distribution of the committed robots (dark violet bars) is very wide, indicating that both resources are selected from time to time as they present advantages and disadvantages. For $P_I = 0.01$, resource quality seems to matter, as the distribution is shifted towards the exploitation of resource B . Higher values of the cross-inhibition probability lead to a larger number of runs ending with a balanced exploitation of both resources, with 50% to 75% of the committed agents exploiting resource A and the remaining ones exploiting resource B . Nevertheless, some runs end up with full commitment for the high quality resource (B).

Overall, these results confirm that the implemented strategy for exploration and balanced exploitation results in expected distributions of robots among the available resources, giving preference to the most profitable one by allocating more robots on it. To evaluate the efficiency of the system, we compare the achieved allocation and retrieval rate with the model empirically obtained in Sect. 4.2. Figure 7 shows the heatmaps representing the ideal efficiency for the case of two resources with different distances ($d_A = 6$ m, $d_B = 8$ m), and for all possible combination of resource qualities ($r_A, r_B \in \{0.01, 0.03, 0.05, 0.1\}$ items/s). On top of the heatmaps, we show a scatter plot corresponding to the results from 100 experimental runs obtained with $P_I = 0.02$ (for additional results in other conditions, see the supplementary material). Each point corresponds to the final allocation of robots to committed or uncommitted populations, and the color corresponds to the experimentally observed efficiency, using the same color coding as for the heatmap. It is possible to note that the scatter plot generally matches the areas where efficiency is high or maximal, especially when the closer resource is not the one with the highest quality ($r_A \neq 0.1$ items/s). Indeed, when A is also a high-quality resource, the distribution is strongly biased towards its exploitation, because the resource is discovered earlier and can sustain a large number of robots. Allocating other robots to the second resource is therefore less probable. When B is the most profitable resource, the distribution

is biased towards its exploitation and is more balanced. Also note that the efficiency in the exploitation is matched between model and experiments, as the colors of the points in the scatter plot closely correspond to the underlying heatmap, hence confirming the suitability of the model we have introduced to evaluate the system efficiency.

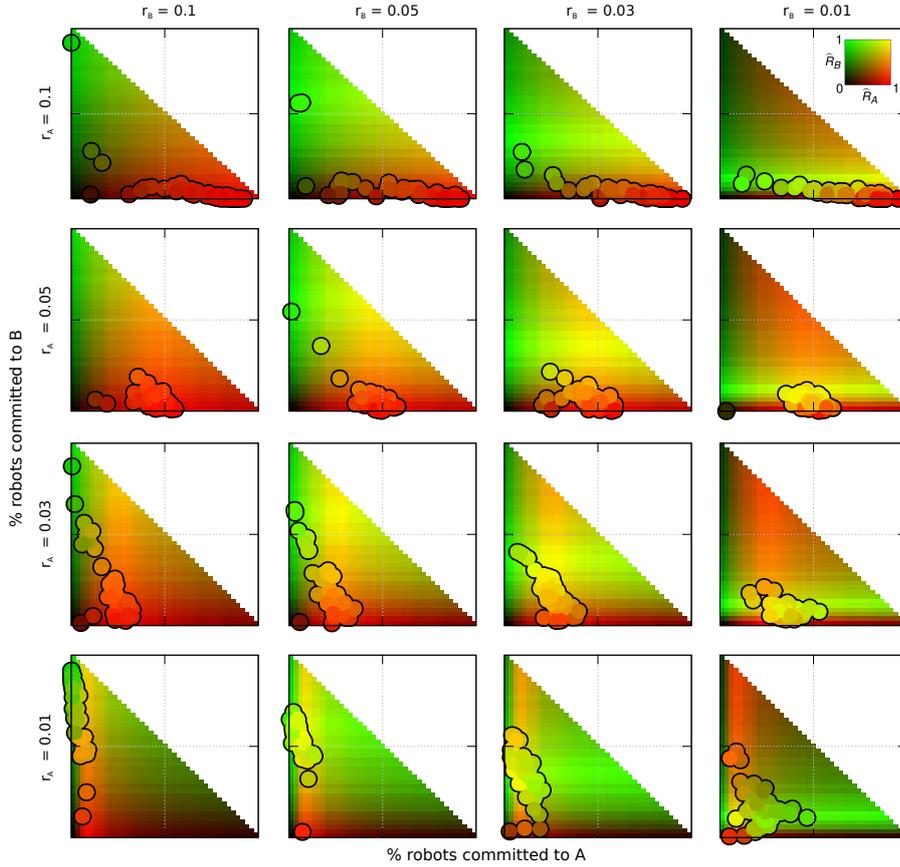


Fig. 7 Efficiency in the exploitation of two resources, placed at different distances ($d_A = 6$ m, $d_B = 8$ m) and for all combination of qualities ($r_A, r_B \in \{0.01, 0.03, 0.05, 0.1\}$ items/s). Each panel shows a ternary plot where the underlying heatmap represents the theoretical efficiency in exploitation following the model introduced in Sect. 4.2, while the scatter plot represents the results of 100 experiments, each point indicating the final allocation of robots to different resources and the observed efficiency in exploitation, represented with the same colour-coding of the heatmap (see Sect. 4.2 and Fig. 4 for details). For additional results in other experimental conditions, see the figures in the electronic supplementary material.

5 Discussion and conclusions

In this work, we have implemented a strategy for exploration and balanced exploitation of renewable resources inspired by the honeybee value-sensitive decision making abilities. We have performed a large-scale simulation analysis to identify the effects of the different parameters governing the individual behaviour on the macroscopic, swarm-level dynamics. The results obtained confirm that our approach is suitable to provide the ability to adaptively balance exploitation of resources at the collective level, without requiring individuals to compare the profitability of different resources, and without a central planner with global knowledge of the environmental conditions.

The decentralised approach we adopt naturally leads to generalisations in the number of resources to be considered and different kinds of exploitation dynamics (as also studied in Reina et al., 2017; Miletitch et al., 2013). Knowledge from modelling studies could be integrated in order to provide parameterisations suitable to deal with more complex working conditions, for instance dealing with a large number of resources in parallel (Reina et al., 2017). However, macroscopic models that consider at the same time the dynamics of the swarm and of the renewable resources are not available to date and require an important analytical effort. Work in this direction is already under way, and can provide means to obtain a precise micro-macro link between the robotics implementation and the modelling predictions, as obtained elsewhere for collective decision-making problems (Reina et al., 2015a,b). The study presented here can be considered the first step towards the definition of a decentralised algorithm capable of optimally dealing with complex and dynamic environmental conditions. Thanks to a wide-ranging analysis of the parameter space, we have demonstrated that the macroscopic dynamics correspond to the expectations, opening up the possibility to develop swarm robotics solutions that appropriately balance the exploitation of resources.

Future work aims at a further development of the proposed behaviour to obtain a more robust implementation. For instance, we found that fine tuning the probability P_G of returning to the nest is complex if one wants to deal with a large range of distances, as P_G strongly influences the average distance from the nest covered by robots while searching. To deal with a large range in the expected distances of resources, the exploration ability of robots should be changed, possibly exploiting recent results on the usage of Lévy walks, which are more suitable for searching in open environments in which the encounter of resources is an episodic event (Dimidov et al., 2016; Schroeder et al., 2017). Advancements in the exploration and exploitation abilities can be obtained also by allowing robots to share information widely while they move, instead of limiting interactions within the nest (Gutiérrez et al., 2010). To this end, it is necessary to understand how the mobility pattern of robots influences their network of interactions, and what is the bearing of an heterogeneous interaction network on the macroscopic dynamics. Additionally, it is worth considering the ability of agents to provide a motion bias to neighbours, thereby including in the study reinforced random walks (Perna and

Latty, 2014; Schroeder et al., 2017). The characterisation of the interaction network resulting from the given mobility pattern can be done in terms of degree distribution and other properties relevant from a network theory point of view (Holme and Saramäki, 2012), while the macroscopic analysis of the effects of the interaction topology on the collective outcome needs to take into account heterogeneous mean field approximations (Moretti et al., 2013).

The way in which different resources are taken into account within the swarm is also worth further investigation. In this work, we limited robots to store only one resource location at the time, therefore constraining the space of possible actions. Different experimental and modelling studies include additional mechanisms that better exploit individual knowledge, such as keeping memory of multiple resources and revisiting previously depleted ones (Dornhaus et al., 2006; Bailis et al., 2010; Granovskiy et al., 2012). Thanks to these mechanisms, higher adaptability is possible against variable environmental conditions. In future work, we will study the possibility of preserving information of known resources even when robots are uncommitted, also linking this aspect with the possibility to develop an emergent language used by the robots to refer to each different option. This is useful especially in case the number and position of resources is not known a priori, so that arriving at consensus on a single label for each resource could be useful to let the robot balance exploitation by interacting in terms of labels and associated features (e.g., estimated regeneration rate or profitability). An interesting aspect to study will be the interaction between the exploitation dynamics and the language dynamics (Steels and Belpaeme, 2005; Loreto et al., 2011), which can lead to synergies between the two processes if these are designed in the correct way (e.g., assign a different label only to the most profitable resource, and the same label to those resources that should not be considered by the swarm). Ideally, robots could attach attributes to labels about resource to enrich the information conveyed (e.g., giving an idea of the exploitation effort so much produced), so that they can evolve a language grounded on their experience (Spranger, 2013). The ability to use more informative communication is key for self-organisation in complex and time-varying environments such as the ones considered in the present study, and beyond.

Acknowledgments

Vito Trianni acknowledges support by the European Commission FP7 Programme People: Marie-Curie Actions through the project “DICE, Distributed Cognition Engineering” (grant agreement number 631297). Marco Dorigo acknowledges support from the Belgian F.R.S.-FNRS, of which he is a Research Director.

References

- P. Bailis, R. Nagpal, and J. Werfel. Positional Communication and Private Information in Honeybee Foraging Models. In *Swarm Intelligence*, pages 263–274. Springer Berlin Heidelberg, Berlin, Heidelberg, 2010.
- F. Bartumeus, M. G. E. da Luz, G. M. Viswanathan, and J. Catalan. Animal search strategies: a quantitative random-walk analysis. *Ecology*, 86(11):3078–3087, 2005.
- E. Bonabeau, G. Theraulaz, and J.-L. Deneubourg. Quantitative study of the fixed threshold model for the regulation of division of labour in insect societies. *Proceedings of the Royal Society of London. Series B: Biological Sciences*, 263(1376):1565–1569, 1996.
- M. Bonani, V. Longchamp, S. Magnenat, P. Rétonnaz, D. Burnier, G. Roulet, F. Vaussard, H. Bleuler, and F. Mondada. The marXbot, a miniature mobile robot opening new perspectives for the collective-robotic research. In *Proceedings of the 2010 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 4187–4193. IEEE Press, 2010.
- J. Borenstein and Y. Koren. Real-time obstacle avoidance for fast mobile robots. *IEEE Transactions on Systems, Man, and Cybernetics*, 19(5):1179–1187, 1989.
- E. Castello, T. Yamamoto, F. D. Libera, W. Liu, A. F. T. Winfield, Y. Nakamura, and H. Ishiguro. Adaptive foraging for simulated and real robotic swarms: the dynamical response threshold approach. *Swarm Intelligence*, 10(1):1–31, 2015.
- F. A. A. Cheein and R. Carelli. Agricultural robotics: Unmanned robotic service units in agricultural tasks. *IEEE Industrial Electronics Magazine*, 7(3):48–58, 2013.
- C. Dimidov, G. Oriolo, and V. Trianni. Random walks in swarm robotics: An experiment with kilobots. In M. Dorigo, M. Birattari, X. Li, M. López-Ibáñez, K. Ohkura, C. Pinciroli, and T. Stützle, editors, *Proceedings of the 10th International Conference on Swarm Intelligence (ANTS 2016)*, volume 9882 of *LNCS*, pages 185–196. Springer New York, 2016.
- M. Dorigo, D. Floreano, L. Gambardella, F. Mondada, S. Nolfi, T. Baaboura, M. Birattari, M. Bonani, M. Brambilla, A. Brutschy, D. Burnier, A. Campo, A. L. Christensen, A. Decugniere, G. Di Caro, F. Ducatelle, E. Ferrante, A. Forster, J. Martinez Gonzales, J. Guzzi, V. Longchamp, S. Magnenat, N. Mathews, M. Montes de Oca, R. O’Grady, C. Pinciroli, G. Pini, P. Rétonnaz, J. Roberts, V. Sperati, T. Stirling, A. Stranieri, T. Stützle, V. Trianni, E. Tuci, A. E. Turgut, and F. Vaussard. Swarmanoid: A novel concept for the study of heterogeneous robotic swarms. *IEEE Robotics & Automation Magazine*, 20(4):60–71, 2013.
- A. Dornhaus, F. Klügl, C. Oechslein, F. Puppe, and L. Chittka. Benefits of recruitment in honey bees: effects of ecology and colony size in an individual-based model. *Behavioral Ecology*, 17(3):336–344, 2006.
- F. Ducatelle, G. A. Di Caro, A. Forster, M. Bonani, M. Dorigo, S. Magnenat, F. Mondada, R. O’Grady, C. Pinciroli, P. Rétonnaz, V. Trianni, and L. M.

- Gambardella. Cooperative navigation in robotic swarms. *Swarm Intelligence*, 8(1):1–33, 2014.
- B. Granovskiy, T. Latty, M. Duncan, D. J. T. Sumpter, and M. Beekman. How dancing honey bees keep track of changes: the role of inspector bees. *Behavioral Ecology*, 23(3):588–596, 2012.
- A. Gutiérrez, A. Campo, F. Monasterio-Huelin, L. Magdalena, and M. Dorigo. Collective decision-making based on social odometry. *Neural computing & applications*, 19(6):807–823, 2010.
- J. P. Hecker and M. E. Moses. Beyond pheromones: evolving error-tolerant, flexible, and scalable ant-inspired robot swarms. *Swarm Intelligence*, 9(1):1–28, 2015.
- P. Holme and J. Saramäki. Temporal networks. *Physics Reports*, 519(3):97–125, 2012.
- C. Hui. Carrying capacity, population equilibrium, and environment’s maximal load. *Ecological Modelling*, 192(1-2):317–320, 2006.
- M. J. B. Krieger, J.-B. Billeter, and L. Keller. Ant-like task allocation and recruitment in cooperative robots. *Nature*, 406(6799):992–995, 2000.
- T. H. Labella, M. Dorigo, and J.-L. Deneubourg. Division of labor in a group of robots inspired by ants’ foraging behavior. *ACM Transactions on Autonomous Adaptive Systems*, 1(1):4–25, 2006.
- S. Liemhetcharat, R. Yan, and K. P. Tee. Continuous foraging and information gathering in a multi-agent team. In *Proceedings of the 2015 International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, pages 1325–1333. International Foundation for Autonomous Agents and Multiagent Systems, Richland, SC, 2015.
- W. Liu and A. F. T. Winfield. Modeling and optimization of adaptive foraging in swarm robotic systems. *The International Journal of Robotics Research*, 29(14):1743–1760, 2010.
- W. Liu, A. F. T. Winfield, J. Sa, J. Chen, and L. Dou. Towards energy optimization: Emergent task allocation in a swarm of foraging robots. *Adaptive Behavior*, 15(3):289–305, 2007.
- V. Loreto, A. Baronchelli, A. Mukherjee, A. Puglisi, and F. Tria. Statistical physics of language dynamics. *Journal of Statistical Mechanics: Theory and Experiment*, 2011(04):P04006, 2011.
- R. Miletitch, V. Trianni, A. Campo, and M. Dorigo. Information aggregation mechanisms in social odometry. In *Proceedings of the 20th European Conference on Artificial Life (ECAL 2013)*, pages 102–109. MIT Press, Cambridge, MA, 2013.
- P. Moretti, A. Baronchelli, M. Starnini, and R. Pastor-Satorras. Generalized voter-like models on heterogeneous networks. In A. Mukherjee, M. Choudhury, F. Peruani, N. Ganguly, and B. Mitra, editors, *Dynamics On and Of Complex Networks, Volume 2: Applications to Time-Varying Dynamical Systems*, pages 285–300. Springer New York, 2013.
- R. R. Murphy, S. Tadokoro, D. Nardi, A. Jacoff, P. Fiorini, H. Choset, and A. M. Erkmen. Search and rescue robotics. In *Springer Handbook of Robotics*, pages 1151–1173. Springer, 2008.

- D. Pais, P. M. Hogan, T. Schlegel, N. R. Franks, N. E. Leonard, and J. A. R. Marshall. A mechanism for value-sensitive decision-making. *PLoS ONE*, 8(9):e73216, 2013.
- A. Perna and T. Latty. Animal transportation networks. *Journal of The Royal Society Interface*, 11(100):20140334–20140334, 2014.
- C. Pinciroli, V. Trianni, R. O’Grady, G. Pini, A. Brutschy, M. Brambilla, N. Mathews, E. Ferrante, G. Di Caro, F. Ducatelle, M. Birattari, L. M. Gambardella, and M. Dorigo. ARGoS: a modular, parallel, multi-engine simulator for multi-robot systems. *Swarm Intelligence*, 6(4):271–295, 2012.
- L. Pitonakova, R. Crowder, and S. Bullock. Information flow principles for plasticity in foraging robot swarms. *Swarm Intelligence*, 10(1):33–63, 2016.
- A. Reina, R. Miletitch, M. Dorigo, and V. Trianni. A quantitative micro-macro link for collective decisions: the shortest path discovery/selection example. *Swarm Intelligence*, 9(2-3):75–102, 2015a.
- A. Reina, G. Valentini, C. Fernández-Oto, M. Dorigo, and V. Trianni. A design pattern for decentralised decision making. *PLoS ONE*, 10(10):e0140950–18, 2015b.
- A. Reina, J. A. R. Marshall, V. Trianni, and T. Bose. Model of the best-of-n nest-site selection process in honeybees. *Physical Review E*, 95(5):052411–15, May 2017.
- J. Roberts, T. S. Stirling, J.-C. Zufferey, and D. Floreano. 2.5D infrared range and bearing system for collective robotics. In *Proceedings of the 2009 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 3659–3664. IEEE Press, 2009.
- N. Saleh and L. Chittka. Traplining in bumblebees (*Bombus impatiens*): a foraging strategy’s ontogeny and the importance of spatial reference memory in short-range foraging. *Oecologia*, 151(4):719–730, 2006.
- A. Schroeder, S. Ramakrishnan, M. Kumar, and B. Trease. Efficient spatial coverage by a robot swarm based on an ant foraging model and the lévy distribution. *Swarm Intelligence*, 11(1):39–69, 2017.
- T. D. Seeley, P. K. Visscher, T. Schlegel, P. M. Hogan, N. R. Franks, and J. A. R. Marshall. Stop Signals Provide Cross Inhibition in Collective Decision-Making by Honeybee Swarms. *Science*, 335(6064):108–111, 2012.
- S. J. Simpson, R. M. Sibly, K. P. Lee, S. T. Behmer, and D. Raubenheimer. Optimal foraging when regulating intake of multiple nutrients. *Animal Behaviour*, 68(6):1299–1311, 2004.
- Z. Song and R. T. Vaughan. Sustainable robot foraging: Adaptive fine-grained multi-robot task allocation for maximum sustainable yield of biological resources. In *Proceedings of the 2013 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 3309–3316. IEEE Press, 2013.
- M. Spranger. Evolving grounded spatial language strategies. *Künstliche Intelligenz*, 27(2):97–106, 2013.
- L. Steels and T. Belpaeme. Coordinating perceptually grounded categories through language: A case study for colour. *The Behavioral and brain sciences*, 28(04):1–61, 2005.

-
- V. Trianni and A. Campo. Fundamental collective behaviors in swarm robotics. In J. Kacprzyk and W. Pedrycz, editors, *Springer Handbook of Computational Intelligence*, pages 1377–1394. Springer, 2015.
- V. Trianni and M. Dorigo. Emergent collective decisions in a swarm of robots. In *Proceedings of the 2005 IEEE Swarm Intelligence Symposium (SIS 2005)*, pages 241–248, 2005.
- G. Valentini, E. Ferrante, and M. Dorigo. The best-of-n problem in robot swarms: Formalization, state of the art, and novel perspectives. *Frontiers in Robotics and AI*, 4:1–43, 2017.
- A. F. Winfield. Foraging robots. In *Encyclopedia of Complexity and Systems Science*, pages 3682–3700. Springer New York, 2009.
- K. Yoshida. Achievements in space robotics. *IEEE Robotics & Automation Magazine*, 16(4):20–28, 2009.