

Understanding the Role of Recruitment in Collective Robot Foraging

Lenka Pitonakova^{1,2,3}, Richard Crowder² and Seth Bullock^{1,2}

¹Institute for Complex Systems Simulation and ²School of Electronics and Computer Science

University of Southampton, United Kingdom

³l.pitonakova@soton.ac.uk

Abstract

When is it profitable for robots to forage collectively? Here we compare the ability of swarms of simulated bio-inspired robots to forage either collectively or individually. The conditions under which recruitment (where one robot alerts another to the location of a resource) is profitable are characterised, and explained in terms of the impact of three types of interference between robots (physical, environmental, and informational). Key factors determining swarm performance include resource abundance, the reliability of shared information, time limits on foraging, and the ability of robots to cope with congestion around discovered resources and around the base location. Additional experiments introducing odometry noise indicate that collective foragers are more susceptible to odometry error.

Introduction

Foraging is a well-studied behaviour in both animals and robots. Just as biological species implement foraging strategies that are adapted to their specific niches, it would be desirable to tailor robot foraging strategies to the particular challenges that they face. For example, large animals like wolves, hyenas or lions forage alone when their food is dispersed, while social insects such as ants, bees and termites recruit their nest mates to collectively obtain food from a patch that can be exploited over a number of visits. When should robots forage collectively?

Since collective strategies that rely on communication, co-ordination, interaction, etc., will tend to require more expensive robots and more programming time than independent, individualist robots, it is important for designers to resort to collective strategies only when they significantly improve collective performance.

In this work, the resource gathering performance of a simulated swarm of collective robots that recruit each other to profitable resource locations is compared with that of a swarm of individualists that forage independently. Swarms are evaluated across a range of scenarios that differ in terms of the challenge involved in locating resource deposits, as well as their spatial distribution and variation in terms of volume and quality. By analysing variation in performance

we aim to uncover which environmental properties favour collective foraging approaches in robot swarms.

Background

Most animals, especially carnivores, perform solitary foraging (Gittleman, 1989). For example, forest birds search for insects alone if food is sparse and many insects are needed (Robinson and Holmes, 1982). Frigate birds searching for fish also do so solitarily, despite the fact that they live in colonies. They tend to disperse while searching in order to optimise their probability of success and rarely return to an area where they have previously fed (Weimerskirch et al., 2004). Similarly, hyenas (Holekamp et al., 2012), lionesses (Stander and Albon, 1993) and chimpanzees (Busse, 1978) hunt alone during periods of sparse prey abundance.

By contrast, some creatures forage collectively, recruiting their collaborators to profitable food locations by sharing information. Recruited foraging is most common in social insects and appears either in the form of stigmergy or direct signalling. Stigmergy is utilised by ants (e.g., Beekman, 2001) and termites (Arab et al., 2012) when they lay pheromone trails that lead to food, changing their environment such that it stores useful information that guides the behaviour of recruited conspecifics.

Bees, however, rely on directly influencing their nest mates. Successful foragers return to the nest to perform a waggle dance that communicates both food quality and location (relative to the dance floor) to watching bees (Seeley, 1994; Granovskiy et al., 2012). Flexibility and sophistication is achieved by both scouting for new food sources and re-evaluation of old foraging sites, allowing bee colonies to rapidly adapt to changing flower quality (Seeley, 1992; Biesmeijer and de Vries, 2001)

A number of approaches have been taken to implementing foraging in robot swarms, including random walking robots that forage independently (e.g., Hoff et al., 2010), bucket brigading swarms where each robot is only responsible for its own portion of a foraging area and items travel towards a base location by being passed between robots (e.g., Shell and Mataric, 2006; Lein and Vaughan, 2009), robots that di-

rectly signal to each other about where items can be found (e.g., Rybski et al., 2004; Jevtic and Gazi, 2010; Sarker and Dahl, 2011), and swarms that use stigmergy, where certain robots act as beacons, maintaining an information gradient towards a deposit (e.g., Hoff et al., 2010).

It is notable that robot swarms are often investigated in a single environmental scenario with a specific distribution of target items (typically a uniform random distribution) or in scenarios that are particularly suited to the swarm strategy. An exception to this was the work of Hoff et al. (2010) who compared a group of randomly walking individualists to a swarm that used stigmergy. However, the individualists in these experiments were highly disadvantaged as they had to find the drop off location by random walk every time they acquired a target and thus performed very poorly.

Here we implement an idealised and simplified version of bee-like recruitment in which robots that are near to each other may exchange subjective information about recently visited resource locations. By exploring a range of environments, we aim to discover both the conditions that favour collective foraging and those that do not.

Methods

Simulation environment

Robots of size 10×10 units were placed in a 4000×4000 continuous-space arena with periodic boundaries (i.e., a torus) featuring a centrally located circular base with a diameter of 100 units. One simulated second consisted of 50 update loops (hereafter “updates”) and an experimental run lasted 600 simulated seconds. One robot travelling at top speed in a straight line could traverse the full length of the world in 80 seconds (i.e., travelling at 50 units per second).

A set of foraging scenarios were defined, each characterised by the number, N_D , of 5×5 resource deposits in the environment and their properties. Each deposit, i , comprised V_i units of material, where each unit had a deposit-specific resource quality, $Q_i \in [0, 1]$, such that the total resource in a deposit was equal to $V_i Q_i$. Each robot could extract up to one unit of volume from a resource deposit per visit, reducing the deposit’s volume and extracting Q_i resource to be returned to the robot base. Once the entire volume of a resource deposit had been collected it ceased to exist in the environment and was not replaced.

By default, all deposits were distributed uniformly at random in space with each deposit having equivalent quality ($Q = 0.5$). However, some environments featured non-uniform spatial distributions comprising a number of spatially clustered *groups* of deposits, and, additionally, some environments featured non-uniform quality distributions in the form of a number of higher-quality *patches*. Note that these patches did not influence the spatial distribution of deposits, only their quality.

Where the number of spatially clustered groups, N_G , was greater than zero, each of these groups featured 10 deposits

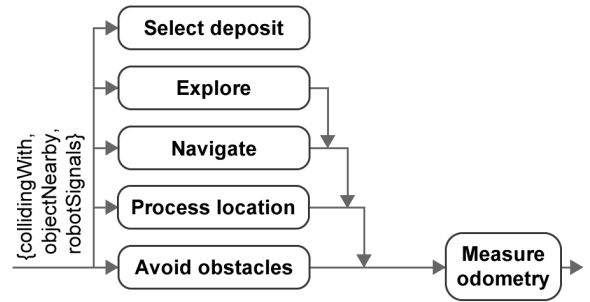


Figure 1: Subsumption architecture of the robot controller.

clustered within a 500-unit radius circular region and no deposits existed outside these groups. Where the number of quality patches, N_P , was greater than zero, each patch was centered on a deposit which was allocated $Q = 1$, and the quality of the deposits around it decayed with distance, d , from the patch centre:

$$Q = \min \left(e^{-d^2 / \beta \sigma_P^2}, 0.5 \right), \quad (1)$$

where patch quality variance, σ_P^2 , was a scenario-specific parameter that controlled the size of the quality patch, and β was a constant factor set to 8×10^6 for all results reported here. In such scenarios, the quality, Q , of deposits outside of patches was scaled down by a constant factor in order to keep the total amount of net resource equal across scenarios. Quality patch centres were randomly chosen deposits with no more than one patch centre per group (when $N_G > 0$).

Robots

The simulated robots utilised a subsumption architecture (Brooks, 1986), where low-level behavioural modules such as avoiding obstacles could be overridden by motor commands of higher-level behaviours such as navigation towards a target (Figure 1). The individual modules communicated with each other by setting the robot’s state and other internal variables stored in the robot’s memory. A finite-state machine representation of the robot is depicted in Figure 2.

Robots were initially randomly oriented and located within the central home base which featured a beacon that allowed any robot to navigate directly home from any point in the arena (Pini et al., 2013).

Individualist robots (I-Swarm) performed a random walk to search for food, avoiding each other by backing up and turning by a random angle. When a deposit was discovered, one unit of volume was loaded (or the entire deposit if one or less than one unit remained) and returned to the base (navigating using the home beacon). On the basis of subjective odometry calculations, a robot returned to where it estimated the deposit was located if the energy efficiency, E_e , of the deposit was better than that of the worst deposit that it had found so far.

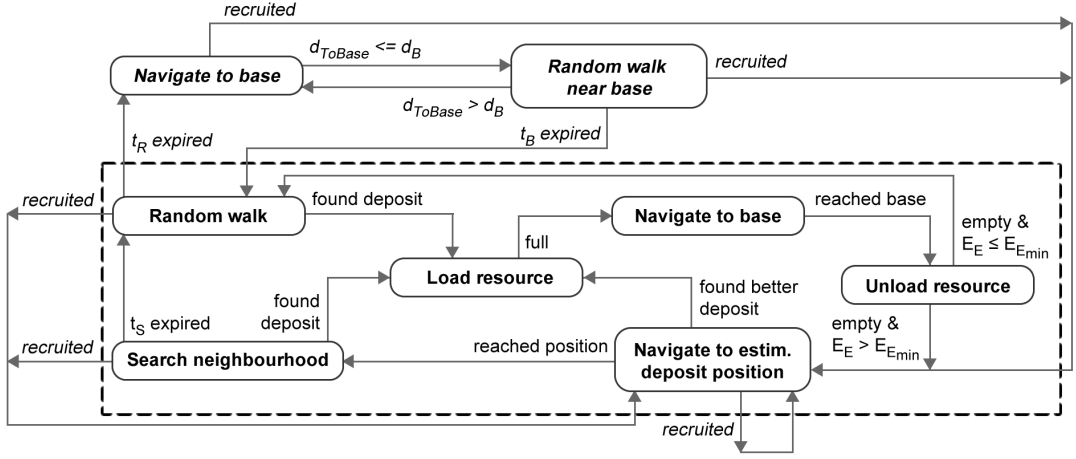


Figure 2: Finite-state machine representation of the robot controller. B-swarm logic is shown in italics outside the dashed box.

Estimation of energy efficiency was inspired by a model of bees (Seeley, 1994) and took into account the volume, V , and quality, Q , of the deposit, as well as the cost of re-visiting it, calculated in terms of time and energy required to (i) reach it (1 unit of energy per simulation update), (ii) load it (5 seconds at 5 units of energy per update), (iii) return to base (5 units of energy per update) and (iv) unload it (2 seconds at 5 units per update):

$$E_e \sim \mathcal{N}\left(\frac{VQ - C}{C}, \sigma_e^2\right) \quad (2)$$

$$C = \sum_{i=1}^4 T_i * E_i, \quad (3)$$

where $\sigma_e^2 = 0.05$ was the variance of normally distributed estimation noise, and T_i and E_i were the time and energy per unit time required for each of the four foraging subtasks. It was also assumed that a robot could estimate the energy efficiency of a deposit from a maximum distance of 30 units.

Upon reaching a remembered location, neighbourhood search was undertaken for a maximum of $t_S = 10$ seconds, comprising a random walk constrained within a distance of $r_S = 120$ units around the estimated deposit location (Pini et al., 2013). The robot resumed random search if nothing was found, or if it could not reach the remembered location within 120 seconds. If, en route to a remembered location, a robot encountered a new deposit with superior energy efficiency, the superior deposit would be visited instead.

Odometry was used to estimate the location of a remembered deposit based on a robot's subjective impression of its rotational change $\Delta\alpha'$ and speed S' during locomotion. Odometry was affected by Gaussian noise on the actual values of rotation (α) and speed (S), simulating wheel slippage,

etc. (Chong and Kleeman, 1997; Martinelli, 2002):

$$\begin{aligned} \Delta\alpha' &\sim \mathcal{N}(\Delta\alpha, \kappa \times \Delta\alpha), \\ S' &\sim \mathcal{N}(S, \kappa \times S), \end{aligned} \quad (4)$$

where parameter κ scaled the variance of odometry noise.

Collective foragers (B-Swarm) employed the same basic foraging strategy as the I-Swarm, but could also signal to nearby robots (< 60 units) their estimate of a remembered deposit's location and its associated energy efficiency (Jevtic and Gazi, 2010). A robot was recruited to a signalled location if it either did not have the location of a deposit in memory, or if its remembered deposit had inferior energy efficiency. Additionally, a B-Swarm robot returned to the base and roamed within $d_B = 200$ units of its centre for up to $t_B = 120$ seconds if it could not find a deposit during $t_R = 120$ seconds of random walk. In this way, unsuccessful foragers could discover deposit locations from robots returning to base to unload, mimicking to some extent the behaviour of bees (Biesmeijer and de Vries, 2001; Granovskiy et al., 2012).

Results

Three sets of simulation results are reported here. First, we explore the impact of varying both the number of robots and the number of deposits on I-Swarm performance in order to identify values for N_R and N_D that present an appropriate degree of challenge. Second, we compare the performance of I-Swarms and B-swarms over a range of foraging scenarios that differ in key respects. Third we explore the impact of odometry error on swarm foraging performance.

Calibrating N_R and N_D

I-Swarm performance was assessed in two basic scenarios, *Uniform* quality and *Patchy* quality, varying the number of

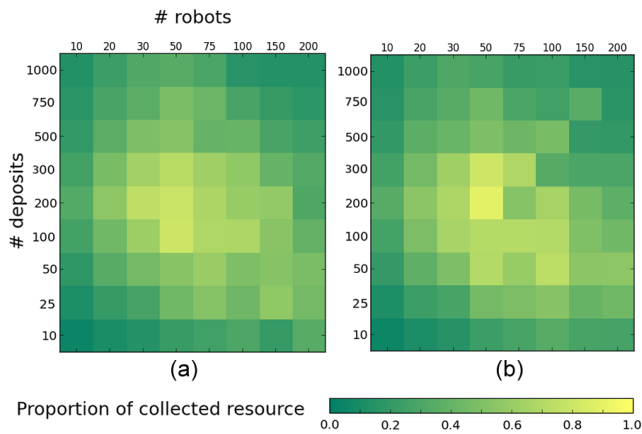


Figure 3: I-Swarm performance as a function of the number of robots and number of deposits in (a) *Uniform* quality and (b) *Patchy* quality scenarios. Each data point is the mean of 10 independent simulations.

robots, N_R , and the number of deposits, N_D (see Figure 3). In both scenarios, $V = 200/N_D$, and deposits were distributed uniformly at random across the entire arena. *Uniform* quality deposits had $Q = 0.5$. The quality of *Patchy* deposits was influenced by $N_P = 30$ randomly located quality patches, each with $\sigma_P^2 = 0.01$.

In both scenarios, performance peaked for an intermediate number of robots and an intermediate number of deposits. With too few robots or too many deposits, not enough foraging trips could be carried out within the duration of a simulation. With too many robots performance deteriorated due to physical interference between foragers near the base. With too few deposits, performance was reduced by the increased challenge involved in discovering deposits far from the base.

In the *Patchy* quality scenario there was some evidence that performance depended on an interaction between N_D and N_R (influencing the degree of interference between robots) but the peak and overall shape of the performance distribution was not significantly altered. Consequently, in order to assess swarms in scenarios that varied in the challenge that they presented, and thereby avoid floor and ceiling effects, a range of values for $N_D \in [10, 100]$ and $N_R \in [10 \dots 100]$ were established and utilised for the remainder of the results presented below.

I-Swarm and B-Swarm Foraging Performance

A set of foraging scenarios were defined, varying in spatial distribution, volume, quality, etc.:

- **Litter** scattered in a field: a uniform random spatial distribution of equal quality deposits, with $N_D = 100$, $N_G = 0$, $N_P = 0$, $V = 2$, $Q = 0.5$.
- **Large Puddles** of rain water: like *Litter*, except deposits are ten times larger, $N_D = 10$, $V = 20$.

- **Rare Minerals**: a uniform random spatial distribution of patchy quality deposits located at least 1500 units from the base, with $N_D = 10$, $N_G = 0$, $N_P = 5$, $\sigma_P^2 = 0$, $V = 20$, $Q = \text{varied}$.
- **Common Stones**: like *Minerals*, but $N_D = 100$, $N_P = 10$, $\sigma_P^2 = 0.01$, $V = 2$, $Q = \text{varied}$.
- **Nectar** from flower clusters: 10 groups of 10 deposits each. Location of group centres was drawn from a uniform random distribution at least 500 units away from the base, with $N_D = 100$, $N_G = 10$, $N_P = 3$, $\sigma_P^2 = 0.01$, $V = 2$, $Q = \text{varied}$.
- **Lost Cargo**: like *Nectar*, but with one group of equal quality deposits where $N_D = 10$, $N_G = 1$, $N_P = 0$, $\sigma_P^2 = 0$, $V = 20$, $Q = 0.5$.

Variation in swarm performance across different scenarios (see Figure 4) can be explained in terms of three types of robot-robot interference:

1. Physical interference caused by one robot obstructing another typically while foraging from the same deposit or returning to the base at the same time.
2. Environmental interference where one robot substantively alters the foraging environment of another by depleting its current target deposit.
3. Informational interference, unique to the B-Swarm, where one robot's behaviour is influenced either positively or negatively by signals from another robot.

I-Swarms slightly but consistently outperformed B-Swarms when collecting *Litter* in swarms of less than 100 robots. This advantage was statistically significant for swarms of size $N_R = 10$ and 20, and also significant across all $N_R < 100$ swarms when their results were considered together (Wilcoxon signed-rank test, $p < 0.01$). The fact that litter deposits were abundant and could be depleted in two visits meant B-Swarms suffered from informational interference that negatively impacted on performance. In theory, *Stones* presented a more significant challenge for both swarms since they were located further from the base and were, at least initially, harder to find. However, the slight advantage that I-Swarms experienced for *Litter* was extinguished in this scenario since B-Swarms could benefit from recruiting robots to the general location of the stones, and were thus able to match I-Swarm performance.

B-Swarms were able to collect more resources than I-Swarms in each of the other scenarios, where deposits were richer in volume but more difficult to discover, i.e., *Puddles*, *Minerals*, *Nectar* and, most strikingly, *Cargo* where all deposits were in one area far from the base. In this scenario, a B-Swarm of 100 robots was even able to collect more resource than an I-Swarm of the same size in an easier, *Nectar* collection task.

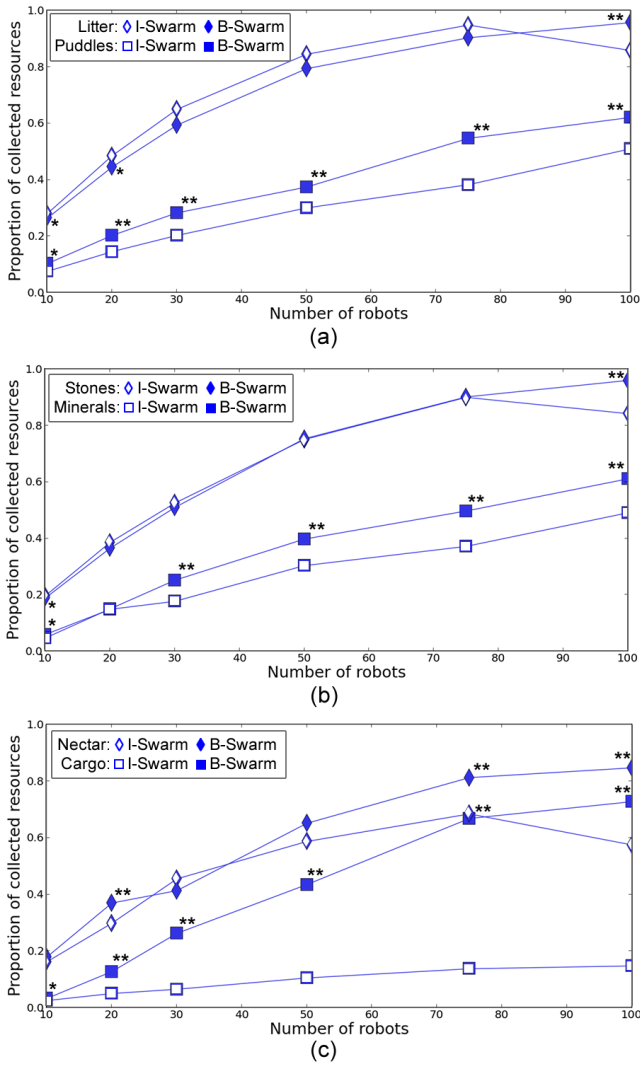


Figure 4: The influence of scenario type on the foraging performance (measured as the median proportion of resource collected over 20 independent simulation runs) of I-Swarms and B-Swarms of different sizes. Statistically significant differences between I-Swarm and B-Swarm performance in the same scenario are indicated with asterisks (Wilcoxon signed-rank test, ** $p < 0.01$, * $p < 0.05$). The average interquartile range was approximately 0.082.

An interesting effect of recruitment was also observed when 100-robot swarms foraged for *Litter*, *Stones* or *Nectar*. I-Swarms of this size suffered a drop in performance resulting from physical interference in the form of congestion around the base. By contrast, recruitment in B-Swarms caused smaller groups of robots to forage from neighbouring locations, meaning that foragers from one group solved the problem of avoiding each other near deposits in parallel with other groups. Such self-organisation led to a kind of spontaneous traffic management and consequent reduction

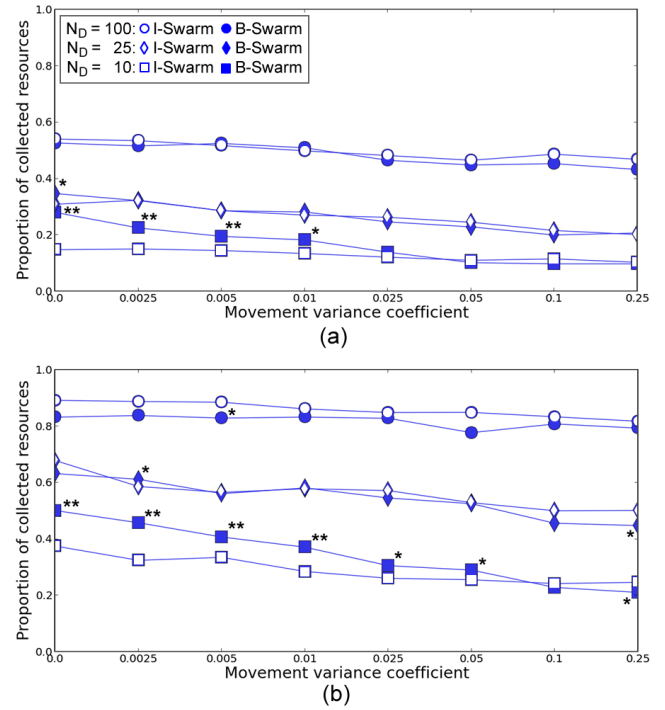


Figure 5: The influence of odometry error (κ) on the performance (measured as the median proportion of resource collected over 20 independent simulation runs) of I-Swarms and B-Swarms of different sizes in the *Variable* scenario: (a) $N_R = 20$, (b) $N_R = 50$. Statistically significant differences between I-Swarm and B-Swarm performance in the same scenario are indicated with asterisks (Wilcoxon signed-rank test, ** $p < 0.01$, * $p < 0.05$). The average interquartile range was approximately 0.072.

of near-base congestion.

On the other hand, there were cases where recruitment did not improve foraging performance. *Mineral* deposits, for instance, were very hard to find for small swarms. B-Swarms needed to have at least 30 members before it was likely that they could find a deposit early enough to ensure that recruitment enabled them to outperform an I-Swarm. Moreover, while B-Swarms typically outperformed I-Swarms when foraging for *Nectar*, this did not hold for swarms of size $N_R = 30$ and 50. Under these conditions, B-Swarms were not large enough to discover and exploit multiple groups of flowers in parallel, but were large enough to rapidly exhaust a single group of flowers by recruitment, ensuring that B-Swarms spent significant time returning to exhausted groups, and proportionately more time randomly searching rather than benefitting from recruitment. This unfortunate interplay between rapid exploitation of a group and insufficient exploration caused B-Swarms of 30 and 50 robots to perform similarly to I-Swarms of the same size.

Odometry error

Error in a robot’s position estimation is an important influence on foraging performance where robots return to remembered deposits or are recruited to new ones. In this section, the value of the κ parameter (Equation 4), previously set to 0.005, is varied in order to explore the effect of odometry noise in the *Patchy* scenario (see Figure 5).

Confirming the previously presented results with low odometry noise, the B-Swarm outperformed the I-Swarm when $N_D = 10$, and performed similarly or slightly worse than the I-Swarm when the number of deposits increased. Increasing κ did not affect the performance of either swarm when deposits were very easy to find ($N_D = 100$) and only slightly degraded the performance when $N_D = 25$.

The impact of odometry error was most significant when 10 deposits were placed in the environment. While collective foragers could harvest significantly more resources than individualists when odometry error was low, the relative advantage of B-Swarms degraded as κ increased until it was extinguished or even slightly reversed. This was an expected result, as the information about location of deposits that guided the B-Swarm robots became more erroneous with high κ , while individualist errors did not propagate to other robots in the I-Swarm, the performance of which was generally less affected by κ .

Interestingly, the impact of odometry noise on B-Swarm performance was stronger for 50-robot swarms than 20-robot swarms. While the performance of a 20-robot B-Swarm rarely fell below that of an equivalent I-Swarm even for high levels of odometry noise, 50-robot B-Swarms tended to perform worse than equivalent I-Swarms for moderate or high levels of odometry noise.

Even though robots in a 50-robot B-Swarm tended to recruit each other at a similar rate, the difficulty of finding advertised deposits when κ was high resulted in increased recruitment to deposits of low quality. For such swarms, the proportion of deposits within a higher-quality patch visited due to recruitment was 21% when $\kappa = 0.01$, but only 13% when $\kappa = 0.1$. Moreover, the distortion of information due to higher odometry error resulted in a smaller total amount of deposit collections (average of 76.9 visits per run when $\kappa = 0.01$, but only 47.1 when $\kappa=0.1$).

In contrast, while increased odometry error also negatively affected the amount of times deposits were visited by a B-Swarm of 20 robots (average of 32.5 times per run when $\kappa=0.01$; 20.6 times when $\kappa=0.1$), the smaller swarm was able to switch to more individualist behaviour when robots failed to locate advertised deposits due to poor odometry: 52% of visited deposits were found by random walk and 33% by recruitment when $\kappa=0.01$, while 63% were found by random walk and only 23% by recruitment when $\kappa=0.1$.

On the other hand, in the larger 50-robot B-Swarm it was harder to behave like an individualist when odometry failed (34% of deposit visits were due to recruitment when $\kappa =$

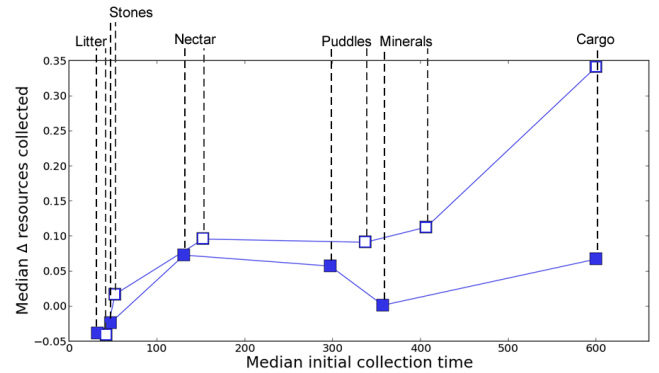


Figure 6: The relationship between the median time taken to collect the first resource and the difference between total resource collected by a B-Swarm and an I-Swarm in various scenarios. Data is shown for 20-robot swarms (solid symbols) and 50-robot swarms (open symbols).

0.01, with a very similar 31% due to recruitment when $\kappa = 0.1$) since a bigger group generated more ultimately fruitless recruitment. If at least one searching robot eventually found a deposit, this information propagated through the swarm, causing the robots to remain near the location and try to access the resource, generating wasteful congestion instead of allowing the resumption of potentially more useful random walk exploration.

Discussion

It is often difficult to intuitively predict the aggregate results or systemic properties of swarm behaviour because of the non-trivial nature of inter-robot interactions and interactions between robots and their environment. However, the analysis presented in this paper shows that it is possible to understand swarms by studying how they behave in various environments. Here we offer a set of generalised principles extracted from the experiments.

When To Forage Collectively

1. When resource is hard to find: B-Swarms performed better than I-Swarms in environments where deposits were harder to find but returned more resource. In particular, the advantage of recruitment was related to the difficulty of discovering deposits (Figure 6), but not to the length of a trip from a deposit to the base, at least within the scale of the scenarios simulated here. The B-Swarm advantage was more significant in larger swarms that covered the area faster and thus had a higher probability of discovering deposits. A small B-Swarm of 20 robots had difficulties, especially where deposit discoveries were rare and represented a small fraction of the total resource. These results suggest that communicating robots are useful for tasks like extracting rare minerals or retrieving lost cargo, but not for picking up litter from streets.




Figure 7: Projected relative gain from using B-Swarm rather than I-Swarm to collect *Nectar*, where the cost of an I-Swarm robot is £300, the gain from collecting all resources is £100,000, and the cost of a B-Swarm robot varies.

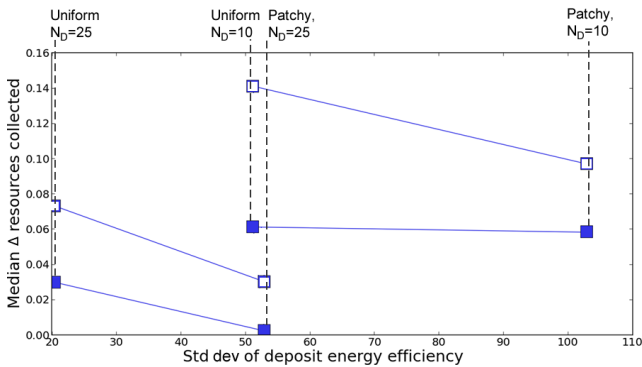


Figure 8: The relationship between the standard deviation of deposit energy efficiency and the performance difference between B-Swarms and I-Swarms. Data is shown for 20-robot swarms (solid symbols) and 50-robot swarms (open symbols). The B-Swarm advantage did not increase in the *Patchy* variant of each *Uniform* scenario.

2. When congestion is a problem: The performance of large swarms of individualists decreased rapidly as congestion around the base prevented robots from foraging. In contrast, recruited groups of B-Swarm robots could solve congestion problems near deposits in parallel with each other, increasing the amount of robots that could operate effectively near the base at the same time. It is notable that this B-Swarm feature would become less effective if the swarm size increased significantly, as the recruited groups would become larger and there would be more of them.

When Not To Forage Collectively

1. When resource is abundant: I-Swarms could slightly outperform B-Swarms when items were abundant and easy to find. Not only was recruitment in these environments unnecessary, but a combination of environmental and informational interference caused advertised locations to become depleted before a recruit could reach them, ensuring that recruitment incurred a performance cost.

2. When reliability of information is low: A decrease in the reliability of information about deposit locations due to odometry error caused the foraging ability of collective foragers to degrade significantly. In these cases, I-Swarm performance may be more reliable, although still generally poor.

3. When deposits are very hard to find quickly: It was very difficult for small swarms of both types to find *Minerals* and *Cargo* where deposits were concentrated in a few distant locations. Despite the additional capabilities of the B-Swarm, once deposits were discovered there was not enough time to recruit others and exploit the location information.

4. Borderline cases: In the real world, the cost of a robot swarm must be reasonable when assessed against the value generated by the foraging task. The additional investment required in order to equip robots with communication transmitters and receivers, and program them with the extended B-Swarm logic should thus be taken into account. For example, imagine that an I-Swarm robot costs £300 and that collecting all resources from the environment is worth £100,000. If the additional cost of a B-Swarm robot is £50, the effect of such an investment is not very dramatic (Figure 7). However, as communicating robots become more expensive, they are less able to return the extra investment.

Further Remarks

As currently implemented, B-Swarms were not better equipped than I-Swarms to selectively forage from more energy efficient deposits (Figure 8). Since B-Swarm robots were recruited to any deposit with an energy efficiency higher than the lowest known E_e , a large proportion of the swarm would collect from mediocre locations instead of exploring the environment. However, it was observed that unemployed B-Swarm robots could target higher E_e deposits as a consequence of receiving information from a large number of returning foragers. It is thus possible that the ability to preferentially exploit deposits with high E_e could be achieved by giving the robots a memory of previously encountered resources and an adaptive E_e threshold that would filter social information, as is the case in bees (Seeley, 1994). It would also be possible to implement a region of the robot base that acts as a dance floor in a beehive (Seeley, 1994; Wray et al., 2011), where robots would meet and compare advertised deposits in a more centralised fashion.

Robots designated for exploration would potentially also improve the performance of the swarm. It was shown that active scouting followed by rapid recruitment are very important factors in making bees so successful in dynamic environments (Granovskiy et al., 2012), compared to other social insects like ants (Ribeiro et al., 2009). Individual effects of the bee-inspired behaviours need to be explored in future experiments in order to establish their importance.

Conclusion

The ability of a simulated robotic swarm of individualist foragers and a swarm of collective foragers to harvest resources from various environments was compared. It was shown that recruitment is useful when resources are hard to find and are either clustered spatially or reward multiple visits, as the robots can spend relatively more time exploiting known deposits than exploring for new ones. On the other hand, the additional cost associated with building communicating robots may not be justified when there are many locations from which to gather resource or when the return expected from repeated visits to deposits is low. Furthermore, robots that rely on communication are more susceptible to errors when positional information loses reliability as may occur through odometry error.

Despite the tempting intuition that better or more expensive swarms will deliver better results, the question of whether to equip robots with the ability to communicate is not a trivial one. Just as natural evolution shapes species to become as simple but at the same time as effective as possible within their biological niche, engineered robots can benefit from our ability to step back and consider the nature of the tasks that they face.

Acknowledgements

This work was supported by an EPSRC Doctoral Training Centre grant (EP/G03690X/1).

Source code: <http://lenkaspace.net/lab/swarmSystems/alife14>

References

- Arab, A., Carollo Blanco, Y., and Costa-Leonardo, A. M. (2012). Dynamics of foraging and recruitment behavior in the Asian subterranean termite *Coptotermes gestroi* (Rhinotermitidae). *Psyche*, 2012. doi:10.1155/2012/806782.
- Beekman, M. (2001). Phase transition between disordered and ordered foraging in Pharaohs ants. *P. Natl. Acad. Sci. USA*, 98(17):9703–9706.
- Biesmeijer, J. C. and de Vries, H. (2001). Exploration and exploitation of food sources by social insect colonies: A revision of the scout-recruit concept. *Behav. Ecol. Sociobiol.*, 49(2):89–99.
- Brooks, R. (1986). A robust layered control system for a mobile robot. *IEEE J. Robot. Autom.*, 2(1):114–23.
- Busse, C. D. (1978). Do chimpanzees hunt cooperatively? *Am. Nat.*, 112(986):767–770.
- Chong, K. S. and Kleeman, L. (1997). Accurate odometry and error modelling for a mobile robot. In *IEEE Int. Conf. Robotics & Automation*, pages 2783–2788.
- Gittleman, J. L. (1989). Carnivore group living: Comparative trends. In Gittleman, J. L., editor, *Carnivore Behavior, Ecology and Evolution*, pages 183–207. Cornell University Press.
- Granovskiy, B., Latty, T., Duncan, M., Sumpter, D. J. T., and Beekman, M. (2012). How dancing honey bees keep track of changes: The role of inspector bees. *Behav. Ecol.*, 23(3):588–596.
- Hoff, N., Sagoff, A., Wood, R. J., and Nagpal, R. (2010). Two foraging algorithms for robot swarms using only local communication. In *IEEE Int. Conf. Robotics & Biomimetics, ROBOT 2010*, pages 123–130. IEEE Press.
- Holekamp, K. E., Smith, J. E., Strelhoff, C. C., Van Horn, R. C., and Watts, H. E. (2012). Society, demography and genetic structure in the spotted hyena. *Mol. Ecol.*, 21(3):613–32.
- Jevtic, A. and Gazi, P. (2010). Building a swarm of robotic bees. In *World Automation Congress*, pages 1–6. IEEE Press.
- Lein, A. and Vaughan, R. T. (2009). Adapting to non-uniform resource distributions in robotic swarm foraging through work-site relocation. In *IEEE/RSJ Int. Conf. Intelligent Robots & Systems*, pages 601–606. IEEE Press.
- Martinelli, A. (2002). The odometry error of a mobile robot with a synchronous drive system. *IEEE T. Robot. Autom.*, 18(3):399–405.
- Pini, G., Brutschy, A., Pinciroli, C., Dorigo, M., and Birattari, M. (2013). Autonomous task partitioning in robot foraging: An approach based on cost estimation. *Adapt. Behav.*, 21(2):118–136.
- Ribeiro, P. L., Helene, A. F., Xavier, G., Navas, C., and Ribeiro, F. L. (2009). Ants can learn to forage on one-way trails. *PLoS One*, 4(4):e5024.
- Robinson, S. K. and Holmes, R. T. (1982). Foraging behaviour of forest birds: The relationships among search tactics, diet and habitat structure. *Ecology*, 63(6):1918–1931.
- Rybski, P. E., Larson, A., Veeraraghavan, H., Lapoint, M., and Gini, M. (2004). Communication strategies in multi-robot search and retrieval: Experiences with MinDART. In Alami, R., editor, *Proc. 7th Int. Symp. Dist. Auton. Robotic Systems*, pages 301–310.
- Sarker, M. O. F. and Dahl, T. (2011). Bio-inspired communication for self-regulated multi-robot systems. In Yasuda, T., editor, *Multi-Robot Systems, Trends and Development*, pages 367–392. InTech.
- Seeley, T. D. (1992). The tremble dance of the honey bee: Message and meanings. *Behav. Ecol. Sociobiol.*, 31(6):375–383.
- Seeley, T. D. (1994). Honey bee foragers as sensory units of their colonies. *Behav. Ecol. Sociobiol.*, 34(1):51–62.
- Shell, D. A. and Mataric, M. J. (2006). On foraging strategies for large-scale multi-robot systems. In *IEEE/RSJ Int. Conf. Intelligent Robots & Systems*, pages 2717–2723. IEEE Press.
- Stander, P. E. and Albon, S. D. (1993). Hunting success of lions in a semi-arid environment. *Symp. Zool. S.*, 65:127–143.
- Weimerskirch, H., Le Corre, M., Jaquemet, S., Potier, M., and Marsac, F. (2004). Foraging strategy of a top predator in tropical waters: Great frigatebirds in the Mozambique Channel. *Marine Ecology Progress Series*, 275:297–308.
- Wray, M. K., Klein, B. A., and Seeley, T. D. (2011). Honey bees use social information in waggle dances more fully when foraging errors are more costly. *Behav. Ecol.*, 23(1):125–131.