OPEN ACCESS

University of BRISTOL

Peer reviewed version

Link to published version (if available):
10.1007/978-3-030-00533-7_22

Link to publication record in Explore Bristol Research
PDF-document

University of Bristol - Explore Bristol Research
General rights

The importance of information flow regulation in preferentially foraging robot swarms

Lenka Pitonakova[1], Richard Crowder[2], and Seth Bullock[3]

[1] Dep. of Computer Science, Fac. of Engineering, Univ. of Bristol, Bristol, UK
`contact@lenkaspace.net`
[2] Dep. of Electronics and Computer Science, Fac. of Physical and Applied Sciences,
Univ. of Southampton, Southampton, UK `rmc@ecs.soton.ac.uk`
[3] Dep. of Computer Science, Fac. of Engineering, Univ. of Bristol, Bristol, UK
`seth.bullock@bristol.ac.uk`

Abstract. Instead of committing to the first source of reward that it discovers, an agent engaged in "preferential foraging" continues to choose between different reward sources in order to maximise its foraging efficiency. In this paper, the effect of preferential source selection on the performance of robot swarms with different recruitment strategies is studied. The swarms are tasked with foraging from multiple sources in dynamic environments where worksite locations change periodically and thus need to be re-discovered. Analysis indicates that preferential foraging leads to a more even exploitation of resources and a more efficient exploration of the environment provided that information flow among robots, that results from recruitment, is regulated. On the other hand, preferential selection acts as a strong positive feedback mechanism for favouring the most popular reward source when robots exchange information rapidly in a small designated area, preventing the swarm from foraging efficiently and from responding to changes.

1 Introduction

Instead of committing to the first source of reward that it discovers, an agent engaged in "preferential foraging" continues to choose between different reward sources in order to maximise its foraging efficiency [13]. This foraging behaviour appears in nature at the level of individual creatures, for example in fish [13] and birds [5], as well as at the collective level, for instance in honey bee [2] and ant [28] colonies. While numerous studies have shown that preferential foraging is advantageous for animals, the conditions under which these advantages transfer to biologically-inspired robot swarms are currently unclear.

During robot swarm foraging, individuals are required to search an unknown environment for *worksites* that contain reward. The robots may perform work directly at the worksite locations during *general foraging* (e.g., [16, 24]) or transport resource collected at worksites to a designated deposition area in *central-place foraging* (e.g., [14, 4]). Robots may communicate information about worksites, such as their locations, to other members of the swarm in order to facilitate

faster exploitation of the environment (e.g., [3, 23]). In this paper, the effect of robots preferring worksites with higher utilities, i.e., higher reward returns, is explored in general and central-place foraging tasks, where worksite locations are not known in advance and change over time. The foraging tasks presented here are a paradigm for a number of real-world robot swarm tasks, such as package delivery, environment sampling and resource collection.

Two robot swarms with different communication strategies are studied. In Broadcaster swarms, individuals advertise worksite information to nearby robots while they are near worksites. In Bee swarms, robots exchange information in a designated area, that they return to periodically. Our previous work [20, 22] suggests that Broadcaster swarms outperform Bee swarms in many foraging environments due to their ability to respond to environmental changes faster, but that Bee swarms are more suitable during central-place foraging in environments with a low worksite density, since their recruitment strategy allows the robots to share information with each other relatively quickly. Here we show that the rapid spread of information through Bee swarms damages their performance when individuals preferentially choose worksites with higher utility, since most of the swarm may tend to concentrate on a single worksite, increasing the negative effects of congestion. On the other hand, Broadcasters form small sub-groups that can use preferential foraging to choose non-congested worksites, because the information spread in these swarms is regulated by the limited communication range of robots.

2 Methods

2.1 Simulated environment

A continuous-space experimental arena, identical to that in [22] and containing a central circular base surrounded by circular worksites, was created in the ARGoS simulator [18]. An experimental environment was characterised by the number of worksites, N_W, and worksite distance from the base, $D \in \{5, 9, 13, 17\}$ m. There were two types of environment (Figure 1a,b):

- **HeapN_W:** $N_W \in \{1, 2, 4\}$ high-volume worksites evenly distributed around the base at a distance D from the base edge.
- **Scatter25**: $N_W = 25$ worksites randomly distributed between distance D and $D - 5$m from the base edge.

The total amount of resource in each environment was set to 100 units and the amount of resource units per worksite, $V_W = 100/N_W$. Each worksite had a 0.1 m radius and there was a colour gradient with 1 m radius around it, that the robots could use to "sense" and navigate towards the worksite (Figure 1c). The base had a radius of 3 m and featured a light source above its centre that the robots could use to navigate towards the base.

At the beginning of each experiment, a number, $N_R \in \{10, 25, 50\}$ of robots were placed at random positions and with random orientations in the base.

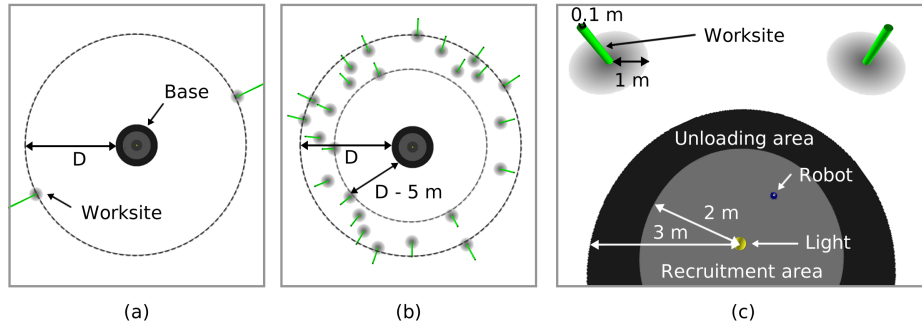Two types of foraging tasks were investigated, as in [22]:

Fig. 1. The (a) Heap2 and (b) Scatter25 environments with worksite distance $D = 13$ m. (c) The base and nearby worksites. Figure reproduced from [22]

- **Consumption**: Worksites represented "jobs" that could be completed at the worksite locations. A robot that was at a worksite gradually depleted its volume, increasing the swarm's total reward at the rate of $1/400$ units per second. Similar tasks were explored in, e.g., [16, 24].
- **Collection**: Worksites represented resource deposits. A robot could collect one unit of resource at a time, after which the resource had to be deposited in the base. Similar tasks were explored in, e.g., [14, 4].

Additionally, each task had two variants, *slow* and *fast*, that represented different degrees of challenge. In each variant, worksite locations changed every T_C seconds and were chosen randomly according to the environment type. For example, in the Heap2 environment with $D = 5$ m, the two worksites were relocated every T_C seconds, remaining 5 m from the edge of the base. Worksite volumes were replenished after each change. The value of T_C, as well as the total simulation time, T, were set as in [22], so that the environment changed 10 times in the slow variant and 20 times in the fast variant and so that a swarm could deplete around 50% of total worksite volume in each slow change interval. For example, $T_C = 45$ min and $T = 7.5$ h for 50-robot swarms.

2.2 Robots

The simulated MarXbots [1] were differentially steered circular robots with a radius of 8.5cm. The robots could communicate with each other using a range and bearing module with a signal range of 5 m. We have previously described the robot model in [20]. There were two types of robot swarm, that we parametrised for the best performance in a series of environments [22]:

- **Broadcaster** (Figure 2a): Robots left the base immediately at the beginning of an experiment to start scouting for worksites. Upon discovering a worksite, a scout started working on it, meaning that it either started depleting its volume (in the Consumption task) or started travelling between the worksite
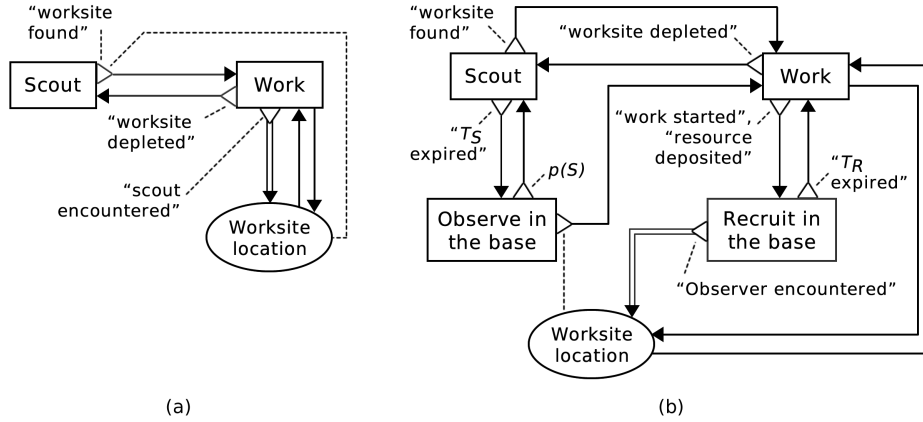
Fig. 2. BDRML [21] representations of the (a) Broadcaster and (b) Bee robot controllers.

and the base to gradually deposit resource (in the Collection task). The robot kept track of worksite location by using odometry. A robot that was located on the gradient surrounding a worksite also broadcasted the worksite location to other robots that were within the communication range. A scout that received the message was recruited to the worksite and started working from it.

– **Bee** (Figure 2b): Robots left the base with probability $p(S) = 10^{-3}$ to start scouting for worksites. In the Consumption task, a robot that discovered a worksite first returned to the recruitment area of the base (Figure 1c) in order to recruit any "observing" robots for $T_R = 120$ s and then resumed working. In the Collection task, a robot recruited after each time that it deposited resource in the base. A scout that could not find any worksites for $T_S = 18$ min returned to the base in order to start observing recruitment signals.

The Broadcaster control strategy was inspired by the behaviour of animals such as sheep [17] and fish [15], where individuals observe each other during foraging and are attracted to locations that others forage from. A similar strategy has also been implemented in robot swarms, e.g. in [3, 8, 30]. The Bee control strategy was inspired by the foraging behaviour of honey bees [27] and has also been implemented in robot swarms, e.g. in [14, 23, 9].

Differential steering sensors and motors of the robots were subject to minor noise, which could result in accumulation of errors in the relative vector to a robot's worksite. Therefore, upon arriving to a supposed worksite location that was empty, a robot performed neighbourhood search around the location that lasted for 5 minutes.

Experiments were performed with *Committed* and *Preferential* swarms. Robots in the Committed swarms remained foraging from the same worksite that they

discovered or were recruited to until the worksite was depleted. Robots in Preferential swarms exchanged both worksite locations and worksite utilities and always preferred to forage from a worksite with a higher utility. In the Consumption task, worksite utility, U_W, was equal to its current volume, V_W:

$$U_W = V_W \tag{1}$$

In the Collection task, the distance, D_W, between a worksite and the base was also considered, so that worksites that were further away from the base had a lower utility:

$$U_W = V_W/D_W \tag{2}$$

A preferentially foraging robot switched to a worksite with a higher utility if it found such a worksite during its journey to its current worksite location. Additionally, in Preferential Broadcaster swarms, a scout could receive recruitment signals from multiple directions and always chose a worksite with the highest advertised U_W. Furthermore, if the distance between worksites was smaller than the communication range of robots, the robots exchanged information about U_W each second while they were working. If worksite A was being depleted faster than worksite B (because more robots were working at A), its U_W decreased faster, meaning that robots from A were eventually recruited to B. In Preferential Bee swarms, robots exchanged information about U_W while they were recruiting in the base. If there were multiple worksites advertised at the same time, all observing robots and recruiters adopted the worksite with the highest advertised utility.

2.3 Terminology and data visualisation

Swarm performance analysis is conducted within the Information-Cost-Reward framework [22]. The framework allows us to identify various costs that the robots incur each second during foraging instead of obtaining reward. The *uncertainty cost*, C_U, is incurred by robots that do not know about worksites. The *displacement cost*, C_D is incurred by robots that know where worksites are located, but are currently not at their worksites, unable to receive reward from them. For example, Bee swarm robots usually incur high amounts of C_D, because they are recruited in the base, i.e., far away from worksites. The displacement cost coefficient, d, represents a ratio between the amount of C_D and the decrease in C_U paid at a given time. When $d = 1$, all robots that know about worksites are displaced from them and no reward is obtained. Intermediate values of $0 < d < 1$ indicate that some robots are displaced and some are receiving reward. Finally, the *misinformation cost*, C_M, is incurred by robots that are away from their worksites and do not know that the worksites have already been depleted by other robots. There is a high potential for a robot to incur C_M during Collection, since it periodically returns to the base to deposit resource.

In the following two sections, swarm performance and ICR metrics are presented in the form of box plots. Each data point represents a median result of 50

independent simulation runs. The surrounding boxes represent the inter-quartile range of the result set, and whiskers represent data in the range of 1.5 times the inter-quartile. Outliers outside this range are shown as plus signs.

3   The performance of Committed swarms

In each foraging task, the performance of the Committed swarms depended on the number of worksites, N_W, and the distance of worksites from the base, D (Figure 3). The performance was generally better in environments with a high worksite density, i.e., when N_W was high, or when D was low. However, note that in the slow Consumption task (Figure 3a), the Bee swarms experienced a high amount of congestion around worksites in Scatter25 environments when $D \leq$ 9 m, causing their performance to be lower than that in Heap4 environments. The nature of their recruitment strategy in this particular task and environment caused many Bee robots to concentrate on a small number worksites, which was disadvantageous when many worksites needed to be found and exploited.

The Broadcasters outperformed the Bee swarms in the Consumption task, because they did not return to the base in order to recruit, which allowed them
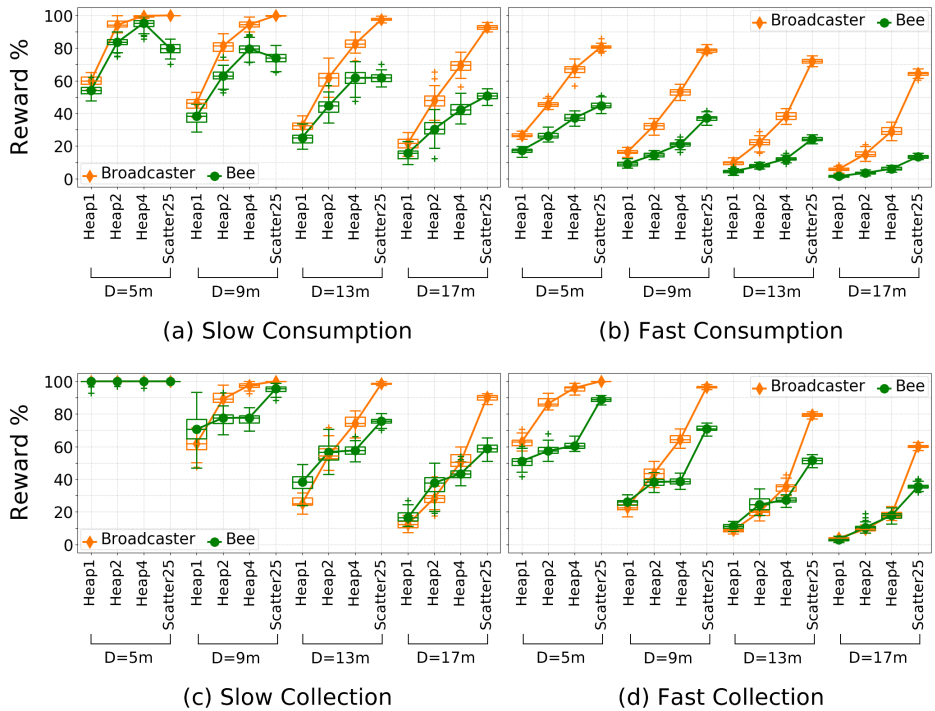

Fig. 3. The amount of reward collected by 50-robot Committed swarms in various tasks and environments.

to spend more time working. However, in the Collection task, the displacement and misinformation costs associated with central-place recruitment were ameliorated by the fact that all robots had to return to the base periodically in order to deposit resource. This allowed the Bee swarms to surpass or match the performance of Broadcasters in many low-density environments, such as Heap1 and Heap2, especially in the slow Collection task. Similar trends in absolute and relative swarm performance were discovered for a number of explored swarm sizes ($N_R \in \{10, 25, 50\}$). However, the largest swarms showed the largest differences in relative performance to each other.

4   The performance of Preferential swarms

The impact of preferential foraging on swarm performance depended on the environment and task type, as well as on the control strategy of robots. In general, swarm performance was affected more strongly when worksites were closer to the base, i.e., when robots could find worksite information faster. Similarly, larger swarms were affected in a larger number of environments, since the robots could receive information from a larger number of recruiters. Broadcasters were able to improve their performance in the Consumption task, while Bee swarms were negatively affected in both tasks (Figure 4).

There were two different ways in which Preferential Broadcasters improved their performance in the Consumption task. In Scatter25 environments with a short D, where it was possible for multiple worksites to be located within the communication range of robots, a recruited robot was able to find worksites with larger volumes, i.e. a larger utility, on its way to the location that it was originally recruited to. This allowed the Preferential Broadcaster swarms to spread their foraging effort across multiple worksites better and prevent congestion, which in many environments decreased their displacement cost coefficient compared to the Committed swarms (Figure 5a,c). This was especially advantageous when the environment changed quickly, i.e., when it was more important to exploit as many worksites as possible in a relatively short amount of time. For example, 50-robot Preferential broadcaster swarms obtained around 12% more reward in the Scatter25 environment when $D = 5$ m (Figure 4b). Secondly, in Heap environments with a small D, recruits often could not reach an advertised worksite due to congestion around it. Robots in congested areas often made a lot of turns while avoiding each other, causing their odometry-based vector to the worksite to become increasingly incorrect due to the cumulative effect of sensory-motor noise. Preferential Broadcasters communicated about the worksite utility and location periodically, meaning that some recruits were sent to incorrect locations. This cleared congestion and allowed the recruits to explore a new area after they could not find the advertised worksite.

On the other hand, in the Collection task, where robots deposited resource in the base and where congestion around worksites was thus cleared periodically, recruitment to incorrect worksite locations decreased the performance of Preferential Broadcasters in Heap environments with a small D (Figure 4c,d). In these
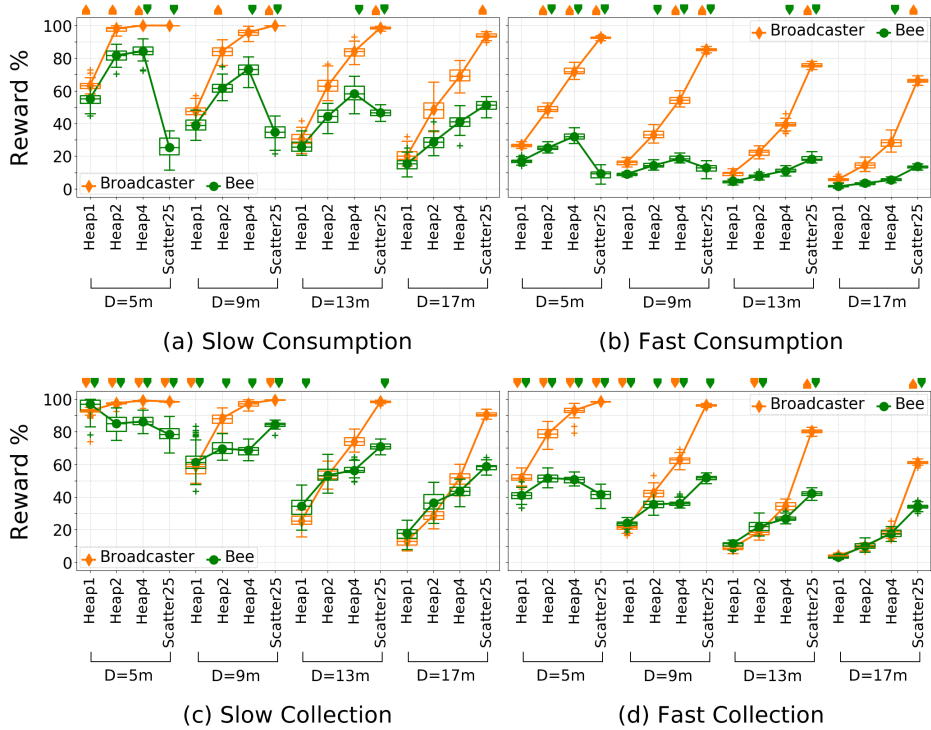
Fig. 4. The amount of reward collected by 50-robot Preferential swarms in various tasks and environments. The up and down arrows above the data points for each environment indicate statistically significant (ANOVA, $p = 0.01$) increase and decrease in performance of the Broadcaster (orange) and Bee (green) Preferential swarms, when compared to the corresponding Committed swarms from Figure 3 in the same environment. When no arrow is shown for a given environment and swarm type, the Preferential and the Committed swarms of that type performed similarly.

cases, it was more advantageous for robots to wait until the path to their worksite became less congested, rather than to travel away and search for new worksites. Secondly, in Scatter25 environments, robots could not share information about worksite utilities with each other as frequently as in the Consumption task, since they spent most of the time travelling between the base and worksites, rather than recruiting near worksites.

Preferential Bee swarms exploited the environment less efficiently than the Committed Bee swarms. Since Bee swarm robots utilised central-place recruitment, the preference for a single worksite with the highest utility often spread to the majority of robots. This increased the amount of congestion around that worksite, preventing the robots from exploiting it, as well as from exploring the rest of the arena. In the Consumption task, the swarm's displacement cost coefficient, d, increased as a result of congestion in environments with multiple work-
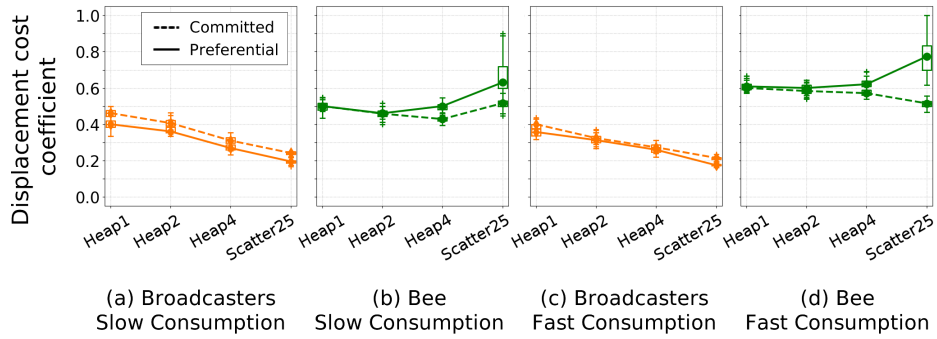
Fig. 5. The displacement cost coefficient, d, of the Committed and Preferential swarms in various tasks and environments with $D = 5$ m.
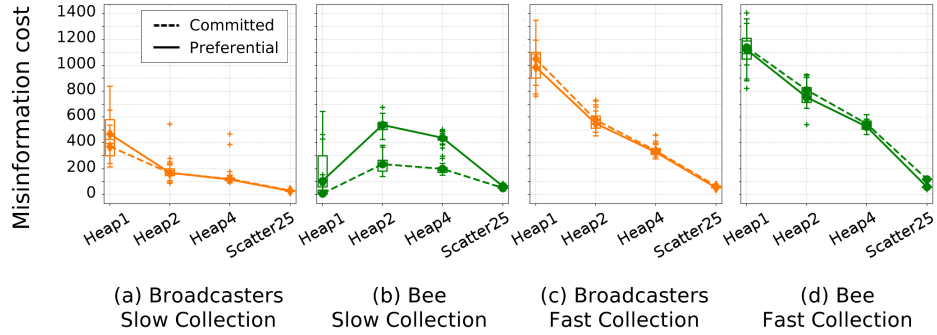

Fig. 6. The misinformation cost, C_M, paid per minute by the Committed and Preferential swarms in various tasks and environments with $D = 5$ m.

sites (Figure 5b,d), while the d of Broadcasters remained similar or decreased (Figure 5a,c). Additionally, in the slow Collection task, where recruitment to a single worksite in Bee swarms was much stronger as the robots exchanged information every time they returned to the base, the amount of misinformation cost that the robots incurred increased (Figure 6b).

As a result of these effects, Preferential foraging decreased the performance of Bee swarms in both tasks, and its impact was significantly stronger than in the Broadcaster swarms (Figure 4). Consequently, the differences between the performance of Preferential Broadcaster and Bee swarms were stronger than between the Committed swarms. Most notably, Preferential Bee swarms performed significantly worse than Broadcasters in environments with a high worksite density due to the increased amount of costs that they incurred.

5 Discussion

While preferential foraging may be beneficial to a foraging swarm, the results presented here indicate that the spread of preference for a selected source needs to be regulated via a negative feedback mechanism, especially when the environment is dynamic and changes over time.

For instance, bee-inspired swarms in [29] were successfully able to distribute their foraging effort between worksites based on worksite utility. Better worksites were advertised in the base for a longer amount of time, while the regulation of information flow was achieved by allowing agents to choose randomly between advertised worksites, preventing all robots from adopting the same choice. In [20], we explored preferentially foraging bee-inspired swarms in dynamic environments where worksite locations remained the same, but where worksite utilities changed over time. The task of the swarms was to collect resource from the worksites into the base and to switch to a worksite with a higher utility when the environment changed. In line with the results presented here, it was shown that bee swarms employing preferential foraging performed worse than other swarms due to their strong commitment to a small portion of worksites. However, presence of negative feedback in the form of utility-dependent worksite abandonment significantly improved the swarm performance.

Studies with other control strategies also showed that maximising information spread within the swarm is not appropriate when a swarm needs to react to a randomly changing environment. In [24], foraging swarms with localised communication outperformed those with a global communication strategy, where all robots were informed about worksite location and urgencies (analogous to our worksite utilities). Localised communication prevented robots from costly travel to distant worksite locations and from interfering with each other near worksites. In ant-inspired foraging swarms, where robots dropped beacons into the environment in order to form virtual trails to worksites, a disrupted trail could be re-established provided that robots sometimes stopped following social information stored in the beacons and started exploring the environment instead [11]. The importance of maintaining a balance between exploration, exploitation and information sharing has also been demonstrated for other collective tasks, such as area clearing [25] and labour division [12].

It is important to note that other aspects of information exchange, for instance, the granularity of data, also impact foraging performance. In [10], ant-inspired swarms formed virtual pheromone trails using beacons between a base and food sources. It was shown that more efficient foraging routes were formed when the beacons stored simple "hop count" integers rather than fine-grained floating point numbers.

Finally, it is interesting that studies of honey bee foraging, contrary to the results presented here, showed that direct exchange of information about nectar profitability among worker bees *improves* the ability of bee colonies to forage efficiently [2]. Workers sample the nectar that other bees bring into the nest, which helps the colony as a whole to react to rapid changes in nectar quality of different flower patches. However, it is important to point out that bees ex-

hibit preferential nectar source selection in addition to a large repertoire of other self-regulatory and communication behaviours. For example, bees also regularly scout the environment in order to discover new flower patches [26] and they periodically check the profitability of old abandoned patches [6]. The colony also monitors and maintains a healthy nectar intake in order to prevent energy wastage resulting from congestion in the nest [7]. Similarly, ants prefer to follow stronger pheromone trails, but pheromone evaporation [28], and in some ant species, time-dependent decrease of interest of free workers in foraging [19], represent negative feedback mechanisms that regulate the colony's responses.

6   Conclusion

In order for preferential source selection to improve swarm performance, the flow of information about worksite utility should be sufficient, but also regulated through a negative feedback mechanism. Sufficient information flow can occur, for example, in environments with an adequate worksite density, or in swarms where robots can meet each other frequently and thus evaluate different sources of information at the same time. Regulation of information flow can be achieved, e.g., by using the Broadcaster recruitment strategy, where recruitment is limited by the communication range of robots. On the other hand, if many robots are allowed to communicate their preferences at the same time, and when robots adopt the best option with a 100% probability, as was the case in our Bee swarms, the environment is exploited and explored inefficiently, preventing the swarm from responding to changes.

These results point to the importance of studying the context in which a particular behaviour is used in a swarm, especially when nature-inspired behaviours are taken out of their biological context and are applied in an engineered robot control algorithm. A new behaviour, such as preferential foraging, adds a new feedback mechanism that interacts with other feedback mechanisms in unique ways. In the experiments presented here, preferential foraging acted as a strong positive feedback mechanism for the most popular worksite choice when robots exchanged information in a small designated area. On the other hand, the same behaviour facilitated a more even exploitation of resources and a more efficient exploration when recruitment was regulated through negative feedback mechanisms.

References

1. Bonani, M., Longchamp, V., Magnenat S., Philippe, R., Burnier, D., Roulet, G., Vaussard, F., Bleuler, H., Mondada, F.: The MarXbot, a miniature mobile

robot opening new perspectives for the collective-robotic research. In: Proc. 2010 IEEE/RSJ Int. Conf. on Intelligent Robots and Systems (IROS 2010). pp. 4187 – 4193. IEEE, Piscataway, NJ (2010)

2. De Marco, R., Farina, W.M.: Trophallaxis in forager honeybees *Apis mellifera*: Resource uncertainty enhances begging contacts? J. Comp. Physiol. A 189, 125–134 (2003)

3. Ducatelle, F., Di Caro, G.A., Forster, A., Bonani, M., Dorigo, M., Magnenat, S., Mondada, F., O'Grady, R., Pinciroli, C., Retornaz, P., Trianni, V., Gambardella, L.M.: Cooperative navigation in robotic swarms. Swarm Intell. 8(1), 1–33 (2014)

4. Ducatelle, F., Di Caro, G.A., Pinciroli, C., Gambardella, L.M.: Self-organized cooperation between robotic swarms. Swarm Intell. 5(2), 73–96 (2011)

5. Gill, F.B., Wolf, L.L.: Nonrandom foraging by sunbirds in a patchy environment. Ecol. 58(6), 1284–1296 (1997)

6. Granovskiy, B., Latty, T., Duncan, M., Sumpter, D.J.T., Beekman, M.: How dancing honey bees keep track of changes: The role of inspector bees. Behav. Ecol. 23(3), 588–596 (2012)

7. Gregson, A.M., Hart, A.G., Holcombe, M., Ratnieks, F.L.: Partial nectar loads as a cause of multiple nectar transfer in the honey bee (*Apis mellifera*): A simulation model. J. Theor. Biol. 222(1), 1–8 (2003)

8. Gutiérrez, Á., Campo, A., Monasterio-Huelin, F., Magdalena, L., Dorigo, M.: Collective decision-making based on social odometry. Neural Comput. Appl. 19(6), 807–823 (2010)

9. Hecker, J.P., Moses, M.E.: Beyond pheromones: evolving error-tolerant, flexible, and scalable ant-inspired robot swarms. Swarm Intell. 9, 43–70 (2015)

10. Hoff, N., Sagoff, A., Wood, R.J., Nagpal, R.: Two foraging algorithms for robot swarms using only local communication. In: Proceedings of the 2010 IEEE International Conference on Robotics and Biomimetics (ROBIO 2010). pp. 123–130. IEEE, Piscataway, NJ (2010)

11. Hrolenok, B., Luke, S., Sullivan, K., Vo, C.: Collaborative foraging using beacons. In: van der Hoek, W., Kaminka, G.A., Lesperance, Y., Luck, M., Sen, S. (eds.) Proc. 9th Int. Conf. on Autonomous Agents and Multiagent Systems (AAMAS 2010). pp. 1197–1204. IFAAMAS, Richland, SC (2010)

12. Jones, C., Mataric, M.J.: Adaptive division of labor in large-scale minimalist multi-robot systems. In: Proc. 2003 IEEE/RSJ Int. Conf. on Intelligent Robots and Systems (IROS 2003). pp. 1969 – 1974, vol. 2. IEEE, Piscataway, NJ (2003)

13. Krause, J., Godin, J.G.J.: Influence of prey foraging posture on flight behavior and predation risk: Predators take advantage of unwary prey. Behav. Ecol. 7(3), 264–271 (1996)

14. Krieger, M.J.B., Billeter, J.B.: The call of duty: Self-organised task allocation in a population of up to twelve mobile robots. Rob. Auton. Syst. 30(1-2), 65–84 (2000)

15. Lachlan, R., Crooks, L., Laland, K.: Who follows whom? Shoaling preferences and social learning of foraging information in guppies. Anim. Behav. 56(1), 181–90 (1998)

16. Lerman, K., Jones, C., Galstyan, A., Mataric, M.J.: Analysis of dynamic task allocation in multi-robot systems. Int. J. Rob. Res. 25, 225–242 (2006)

17. Michelena, P., Jeanson, R., Deneubourg, J.L., Sibbald, A.M.: Personality and collective decision-making in foraging herbivores. Philos. Trans. R. Soc. Lond. B Biol. Sci. 277(1684), 1093–9 (2010)

18. Pinciroli, C., Trianni, V., O'Grady, R., Pini, G., Brutschy, A., Brambilla, M., Mathews, N., Ferrante, E., Caro, G., Ducatelle, F., Birattari, M., Gambardella, L.M.,

Dorigo, M.: ARGoS: A modular, parallel, multi-engine simulator for multi-robot systems. Swarm Intell. 6(4), 271–295 (2012)

19. Pinter-Wollman, N., Bala, A., Merrell, A., Queirolo, J., Stumpe, M.C., Holmes, S., Gordon, D.M.: Harvester ants use interactions to regulate forager activation and availability. Anim. Behav. 86(1), 197–207 (2013)

20. Pitonakova, L., Crowder, R., Bullock, S.: Information flow principles for plasticity in foraging robot swarms. Swarm Intell. 10(1), 33–63 (2016)

21. Pitonakova, L., Crowder, R., Bullock, S.: Behaviour-Data Relations Modelling Language for multi-robot control algorithms. In: Proc. 2017 IEEE/RSJ Int. Conf. on Intelligent Robots and Systems (IROS 2017). pp. 727–732. IEEE, Piscataway, NJ (2017)

22. Pitonakova, L., Crowder, R., Bullock, S.: The Information-Cost-Reward framework for understanding robot swarm foraging. Swarm Intelligence 12(1), 71–96 (2018)

23. Reina, A., Mileticht, R., Dorigo, M., Trianni, V.: A quantitative micro-macro link for collective decisions: the shortest path discovery/selection example. Swarm Intell. 9(2), 75–102 (2015)

24. Sarker, M.O.F., Dahl, T.S.: Bio-Inspired communication for self-regulated multi-robot systems. In: Yasuda, T. (ed.) Multi-Robot Systems, Trends and Development, pp. 367–392. InTech (2011)

25. Schmickl, T., Crailsheim, K.: Throphallaxis within a robotic swarm: Bio-inspired communication among robots in a swarm. Auton. Robots 25(1), 171–188 (2008)

26. Seeley, T.D.: Honey bee foragers as sensory units of their colonies. Behav. Ecol. Sociobiol. 34(1), 51–62 (1994)

27. Seeley, T.D., Camazine, S., Sneyd, J.: Collective decision-making in honey bees: How colonles choose among nectar sources. Behav. Ecol. Sociobiol. 28, 277–290 (1991)

28. Sumpter, D.J.T., Beekman, M.: From nonlinearity to optimality: Pheromone trail foraging by ants. Anim. Behav. 66(2), 273–280 (2003)

29. Valentini, G., Hamann, H., Dorigo, M.: Self-organized collective decision making: The weighted voter model. In: Proc. 13th Int. Conf. on Autonomous Agents and Multiagent Systems (AAMAS 2014). pp. 45–52. ACM, New York (2014)

30. Wawerla, J., Vaughan, R.T.: A fast and frugal method for team-task allocation in a multi-robot transportation system. In: Proc. 2010 IEEE Int. Conf. on Robotics and Automation (ICRA 2010). pp. 1432–1437. IEEE, Piscataway, NJ (2010)