

Beyond Onboard Sensors in Robotic Swarms

Local Collective Sensing through Situated Communication

Tiago Rodrigues, Miguel Duarte, Sancho Moura Oliveira and Anders Lyhne Christensen

Instituto de Telecomunicações, 1049-001 Lisbon, Portugal

Instituto Universitário de Lisboa (ISCTE-IUL), 1649-026 Lisbon, Portugal

BioMachines Lab, 1649-026 Lisbon, Portugal

Keywords: Swarm Robotics, Evolutionary Robotics, Situated Communication, Local Collective Sensing.

Abstract: The constituent robots in swarm robotics systems are typically equipped with relatively simple, onboard sensors of limited quality and range. When robots have the capacity to communicate with one another, communication has so far been exclusively used for coordination. In this paper, we present a novel approach in which local, situated communication is leveraged to overcome the sensory limitations of the individual robots. In our approach, robots share sensory inputs with neighboring robots, thereby effectively extending each other's sensory capabilities. We evaluate our approach in a series of experiments in which we evolve controllers for robots to capture mobile preys. We compare the performance of (i) swarms that use our approach, (ii) swarms in which robots use only their limited onboard sensors, and (iii) swarms in which robots are equipped with ideal sensors that provide extended sensory capabilities without the need for communication. Our results show that swarms in which local communication is used to extend the sensory capabilities of the individual robots outperform swarms in which only onboard sensors are used. Our results also show that in certain experimental configurations, the performance of swarms using our approach is close to the performance of swarms with ideal sensors.

1 INTRODUCTION

Robots in large-scale decentralized multirobot systems, or *swarm robotics systems*, typically have simple and inexpensive sensors. This design principle allows for the unit cost to be kept low, but limits the sensory capabilities of the individual robots (see, for instance, Correll and Martinoli (2006)).

Many simulation-based studies have disregarded limitations of real sensors (Turgut et al., 2008), used simple communication to facilitate cooperation (Fredslund and Matarić, 2002), or relied on indirect coordination through stigmergy (Beckers et al., 1994). While unrealistic sensors can be used to study certain aspects of biological systems, such as the evolution of particular behaviors observed in nature (Trianni et al., 2003; Duarte et al., 2011), resulting controllers cannot be used on any real robotic systems. Simple means of communication, such as sound and color, on the other hand, are relatively straightforward to implement in real hardware (Floreano et al., 2007), but they are also limited in terms of the amount of information they allow robots to exchange. Examples of bio-inspired ap-

proaches such as quorum sensing in bacteria (Bassler, 1999), trophallaxis (Schmickl and Crailsheim, 2008), and hormone-based communication (Stamatis et al., 2009) have been shown to be simple, yet effective strategies to achieve coordination through communication in multirobot systems.

Robots can alternatively be equipped with more complex, wireless communication hardware that enables direct transmission of binary data. In such scenarios, robots are able to transmit packets with relatively large amounts of information. Yet, in swarm robotic systems, such means of communication are typically only used to broadcast simple information, such as the heading, location or speed of each robot (Cianci et al., 2007) in order to facilitate behaviors such as aggregation (Garnier et al., 2008) and flocking (Turgut et al., 2008).

In this paper, we show how the use of communication can extend the limited sensory capabilities of the constituent robots in a swarm. Our approach relies on local, *situated communication* (Støy, 2001), where the signal that carries information also contains context, namely the relative direction and distance,

from the sender to the receiver. This type of communication can be achieved using widely available and relatively inexpensive equipment, such as the e-puck (Mondada et al., 2009) equipped with the range & bearing board extension (Gutiérrez et al., 2008).

In our approach, each robot in the swarm shares readings from its onboard sensors with neighboring robots. By combining information from multiple sources, it becomes possible to obtain knowledge about the environment that would otherwise be unavailable to any single robot. Robots can effectively use the sensory information received to extend their own sensing capabilities through the implementation of virtual sensors. We call such virtual sensors *collective sensors*.

Our goal is to maintain the desirable properties of natural swarm systems while simultaneously exploiting some of the unique capabilities of machines. In this way, we can combine features such as scalability and robustness due to the exclusive reliance on decentralized control, with robots' capacity for low-latency and high-bandwidth communication in order to overcome limitations of the individual units' onboard sensory hardware.

We demonstrate our approach in a predator-prey task, in which a swarm of robots must locate and consume preys. Each robot has two short-range onboard sensors that can detect preys that the robot is facing, and four collective sensors that are implemented based on sensory information shared by neighboring robots. By knowing the neighboring robots' relative range and bearing, it is possible for a robot that receives the information to estimate a prey's position within its own local frame of reference. The estimated position of the prey is then used to compute the readings of the receiving robot's collective sensors thereby effectively extending the robot's sensory range. We compare the performance of robots using collective sensors with robots that rely exclusively on onboard sensors, and with robots equipped with *ideal sensors*, that is, sensors that allow robots to sense preys directly and independently at ranges equal to the collective sensors. In our experiments, we use evolutionary robotics techniques (Nolfi and Floreano, 2000) to evolve artificial neural network-based controllers in scenarios with up to 20 robots and 10 preys.

2 RELATED WORK

Communication systems in nature have been widely studied by biologists, and have served as inspiration to roboticists. The process of communication in bacteria, known as quorum sensing (Bassler, 1999; Einol-

ghozati et al., 2012), relies in producing signaling molecules that can be perceived by neighbors. In this way, it is possible for individuals to estimate population density based on the concentration of signaling molecules, and to modify their behavior accordingly. The quorum sensing process has also been used in robots: Chandrasekaran and Hougen (2006), for instance, used quorum sensing to give nano-scaled robots, constructed using biological components such as proteins and DNA structures, the ability to communicate and coordinate goal-seeking strategies.

Duarte et al. (2011) studied the emergence of complex macroscopic behaviors observed in colonies of social insects, such as task allocation, communication, and synchronization. In a foraging scenario, robots were given explicit visual communication capabilities through changes in the robots' body color. The authors observed that explicit communication enabled complex behaviors to emerge, and the performance of the swarm was significantly higher than in scenarios in which robots could not communicate.

Inspired by mound-building termites, Werfel et al. (2014) implemented a system in which a group of robots were able to coordinate in a construction task. A set of rules was defined, allowing the robots to incrementally construct a particular structure in a decentralized way. The robots did not communicate, and instead had to rely on stigmergy to coordinate.

Turgut et al. (2008) studied self-organized flocking in a swarm of robots with inter-robot communication. Their robots were equipped with a wireless communication module, which allowed the robots in the swarm to sense the headings of neighboring robots. By taking into account the robots' mean orientation, the swarm was able to achieve a robust flocking behavior. In a related study, Fredslund and Mataric (2002) studied formation tasks in swarm of robots. The authors used robots that were equipped with a panning camera and IR sensors. The robots' sensors allowed them to estimate the orientation and distance to other robots in the formation. No global coordinate system was used, and therefore only relative distances were taken into account. The authors enforced formation sorting through each robots' unique IDs using local sensing and minimal communication.

In our study, we go beyond simple communication of each robot's own parameters, such as heading, distance to other robots, or speed. We process onboard sensory information, such as the estimation of the position of a target, and transmit it to neighboring robots. Collective sensors then use the received information to allow the robot to sense particular environmental features that would otherwise be beyond the range of the robot's onboard sensors.

3 METHODOLOGY

In this study, we explore the potential benefits of sharing sensory information to extend the capabilities of the individual robots in swarm robotics systems. The proposed approach is based on the mutual sharing of readings from onboard sensors between neighboring robots. The shared information is then used to compute readings for collective sensors, which can give the individual robot information that would not be available through its onboard sensors.

In our approach, robots can either share preprocessed information, such as the location of interesting features in the environment, or the raw sensory readings, such as readings from a proximity sensor. The shared information is broadcast to nearby robots using situated communication, where the receiving robot knows the relative location and orientation of the transmitting robot. The location and orientation of the transmitting robot can be included in the messages based on GPS and compass information, or by communication means that implicitly embed such relative position information in the signals transmitted (Gutiérrez et al., 2008).

Our collective sensors calculate the appropriate readings taking into account the robot's own location and orientation. This local sensor fusion can provide robots with either longer range sensing, more accurate sensing, or both. For instance, two or more robots observing an object from two different angles may be able to estimate its volume by combining their sensory inputs, something that would not be possible based on readings from a single robot. In this study, robots exchange information regarding the relative position of preys, effectively extending the sensory range of each robot in the swarm.

4 EXPERIMENTAL SETUP

We evaluate our approach in a predator-prey task where a group of robots (the predators) must locate and consume a number of moving preys. The environment is square-shaped, with a size of 10 m x 10 m, surrounded by walls. The robots start each experiment in the center of the environment, while the preys are placed in random locations sampled from a uniform distribution. A robot consumes a prey by touching it. Whenever a prey is consumed, a new prey is placed randomly in the environment, thereby keeping the number of preys constant.

We use small (10 cm diameter) differential-drive robots, loosely modeled after the e-puck (Mondada et al., 2009). The speed of the robots is limited to

10 cm/s. The set of sensors is composed of (i) two onboard prey sensors with a range of 0.8 m, (ii) four collective prey sensors with a range of 3 m, (iii) four robot sensors with a range of 3 m, and (iv) four wall sensors with a range of 0.5 m. All the sensors have an opening angle of 90° . The collective prey sensors, the robot sensors, and the wall sensors are all distributed evenly around the robot, at the angles 0° , 90° , 180° and 270° , while the two onboard prey sensors are located on the front of the robot at angles of 15° and -15° (see Figure 1). Consequently, the onboard prey sensors overlap by 60° and cover a section of 120° . The fact that the two onboard prey sensors overlap was found to help the robot to locate and pursue preys. The two onboard prey sensors could be implemented on real robots, based on inputs from a camera, for instance, by segmenting the field-of-view of the camera into two overlapping regions.




Figure 1: Location and field of view of the two onboard prey sensors, where PL indicates the area sensed by the left onboard prey sensor and PR the right onboard prey sensor. The onboard prey sensors have a range of 0.8 m and opening angle of 90° . Together they cover a 120° wide section, and overlap by 60° .

Readings for the collective sensors are computed based on estimates received from nearby robots that are detecting a prey with their onboard sensors. The relative position of the prey is calculated taking into account the relative distance and orientations of the two robots, as well as the prey's relative location with respect to the robot that is detecting the prey. If estimates are received from multiple robots, it becomes possible to triangulate the position of the prey. Otherwise, an a priori estimate is used in terms of how close the prey is to the robot that is detecting it. An a priori estimate of 50 cm between a prey and a robot is used, which corresponds to 10 times the radius of the robot. The sharing of information is limited to the range of the local, situated communication technology. In this study, the range of both the collective sensors and of local, situated communication is 3 m. An illustration of the collective sensors can be seen in Figure 2.

The preys are able to move at a speed of 15 cm/s, which is 1.5 times faster than the robots. The preys'




Figure 2: An illustration of how collective sensors work. When robot A senses a prey with its onboard sensors, it processes the sensed information generating a estimate of the prey's position and broadcasts the estimate to nearby robots, in this case robot B. Since robot C is outside robot A's communication range, robot A's estimates do not reach robot C.

sensors consist of (i) two wall sensors in front of the prey with a range of 0.5 m — located at the same positions as the prey sensors of the robots, and (ii) four robot sensors, located around the prey's body, with a range of 0.5 m. The prey remains still whenever no nearby robot is detected. If a robot is detected, but it is not directly in front of the prey, the prey moves forward at full speed. If a prey detects a robot in front of it, the prey randomly turns to either side until it is able to move forward. After a prey escapes from the robots, it remains still until a nearby robot is detected again.

We evolved controllers for the robots to solve the proposed task using a simple generational evolutionary algorithm. Each generation was composed of 100 genomes, and each genome encoded an artificial continuous-time recurrent neural network (Beer and Gallagher, 1992) with a reactive layer of input neurons, one hidden layer with 10 neurons, and one layer of output neurons (see (Rodrigues et al., 2014) for a detailed description of the artificial neural network topology used in this study). The fitness of a genome was sampled 9 times and the mean fitness is used for selection. Each sample lasted 5,000 time steps, which is equivalent to 500 seconds. In each

sample, the number of robots and preys were varied in order to promote the evolution of general behavior, which means that one sample was conducted for each possible combination of number of robots and number of preys. The number of robots varied between 5, 10 and 20, and the number of preys varied between 2, 5 and 10. After all the genomes had been evaluated, an elitist approach was used: the top five genomes were selected to populate the next generation. Each of the top five genomes became the parent of 19 offspring. An offspring was created by applying a Gaussian noise to each gene with a probability of 10%. The 95 mutated offspring and the original five genomes constituted the next generation.

In order to evaluate the controllers, we rewarded robots for moving close to and consuming preys, according to the following equations:

$$F = \frac{N_p + \sum_{i=0}^T B_i}{N_r} \quad (1)$$

$$B_i = \sum_{r=0}^{N_r} \max(PL_r, PR_r) \cdot 10^{-5} \quad (2)$$

where N_p is the total number of preys consumed, T is the total number of time steps, N_r is the number of robots on the environment in each sample and $\max(PL_r, PR_r)$ gives the maximum of the readings of the left and right prey sensor for robot r at each time step. Fitness is divided by the number of robots, N_r , in order to prevent biasing evolution toward local optima in setups with many robots. B_i is a bootstrapping term used to guide evolution toward behaviors that result in robots being close to preys.

We ran experiments for three different setups: (i) the collective sensors setup, that represent our approach, (ii) the onboard sensors setup, where robots do not share any information, and (iii) the ideal sensors setup, where robots have sensors that let them sense preys up to a range of 3 m, which is equal to the range of the collective sensors. We ran 20 evolutionary runs in every setup, each lasting 500 generations. After all evolutionary runs had finished, we conducted a post-evaluation with a total of 900 samples, 100 for each combination of numbers of robots and preys, of the genome that had obtained the highest fitness in each run.

For our experiments we used JBotEvolver (Duarte et al., 2014), an open source, multirobot simulation platform and neuroevolution framework.

5 RESULTS AND DISCUSSION

5.1 Performance

Figure 3 shows the average fitness scores of the highest scoring controllers of the collective, onboard and ideal sensors setups. Each boxplot aggregates the results of 900 post-evaluation samples, from 9 different configurations of number of robots and preys. The results show that the highest-performing controllers evolved in the collective sensors setup outperformed the controllers in the onboard sensors setup, and underperformed the controllers in the ideal sensors setup. The average fitness obtained by the highest-performing controllers in post-evaluation of the collective, onboard and ideal sensors setups corresponds to 0.43 ± 0.24 , 0.12 ± 0.04 and 0.72 ± 0.07 , respectively, and in terms of preys consumed to 5.37, 1.61 and 8.81, respectively (see Table 1).

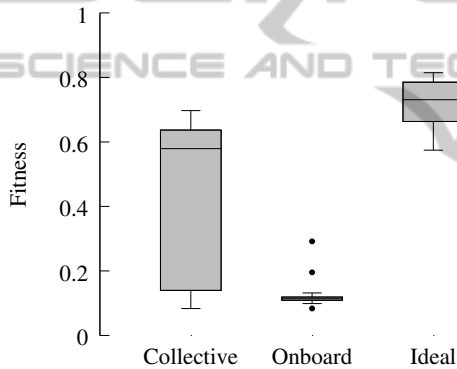


Figure 3: Boxplot of the post-evaluated fitness scores achieved by the highest scoring controllers in 20 evolutionary runs conducted in each of the setups. Each boxplot summarizes results from 900 post-evaluation samples, and comprises observations ranging from the first to the third quartile. The median is indicated by a bar, dividing the box into the upper and lower part. The whiskers extend to the farthest data points that are within 1.5 times the interquartile range, and the dots represent outliers.

Table 1: Mean fitness obtained and number of preys consumed by the highest-performing controllers in post-evaluation of the controllers evolved in the collective, onboard and ideal sensors setups.

	Collective	Onboard	Ideal
Fitness	0.43	0.12	0.72
Preys consumed	5.37	1.61	8.81

When comparing the performance of the controllers evolved in the collective sensors setup with those evolved in the ideal sensors setup, the latter achieved a higher fitness, which can be explained by the fact that the collective sensors need at least one robot detecting a prey with its front prey sensors to be

able to share that prey’s relative position with other nearby robots. In the ideal sensors setup, no communication is necessarily used, since the prey sensors have a range of 3 m instead of 0.8 m, and detect preys in all directions. These differences between the collective and the ideal sensors translate into a mean difference of preys consumed in post-evaluation of 2.68 preys, which corresponds to 30%.

In order to evaluate the robustness, adaptivity and scalability of the solutions evolved, we evaluated the controllers from the highest-performing evolutionary run using collective sensors in an environment where the principal factors — size of the arena, number of preys and robots, were scaled by a factor of five, resulting in an arena of 22.3 m x 22.3 m (500 m²), 50 preys and 100 robots. The evolved controllers were able to disperse well, locate and consume an average of 44.6 preys after post-evaluation, three times the average number of preys consumed by the controller in the original setup (14.8). The number of preys consumed was only three times higher and not five, due to the fact that the average distance from a robot to the wall is longer in the enlarged arena, and robots often need to trap preys in corners or along walls before they can catch them.

5.2 Behavior

The evolved behaviors can be divided into two sub-behaviors: a search behavior and a trap/consume behavior. The preys are faster than the robots, which means that the robots often have to trap a prey before they can catch it. A prey can become trapped if it moves close to a wall or into a corner, and two or more robots are following it closely. Alternatively, three or more robots can trap a prey without the aid of walls by approaching from different directions.

In the collective sensors setup, 15 out of 20 runs evolved the same type of behavior: at the start of a trial, the robots disperse in outward circular motion in order to find preys. The robots then try to chase preys toward the corners or a wall, either in groups or alone. When a prey is consumed or escapes, the robots disperse again to cover a larger area. An example of this behavior can be seen on Figure 4.

The most significant difference found in the behaviors are in the extremes of robot densities, that is, between samples where 20 robots are present and samples where only five robots are present. When the density of robots is high, they tend to disperse evenly and when a prey is seen, they quickly aggregate with nearby robots on the location of the prey. On the other hand, when the density of robots is low, the aggregation near the prey is slower since the robots tend to be

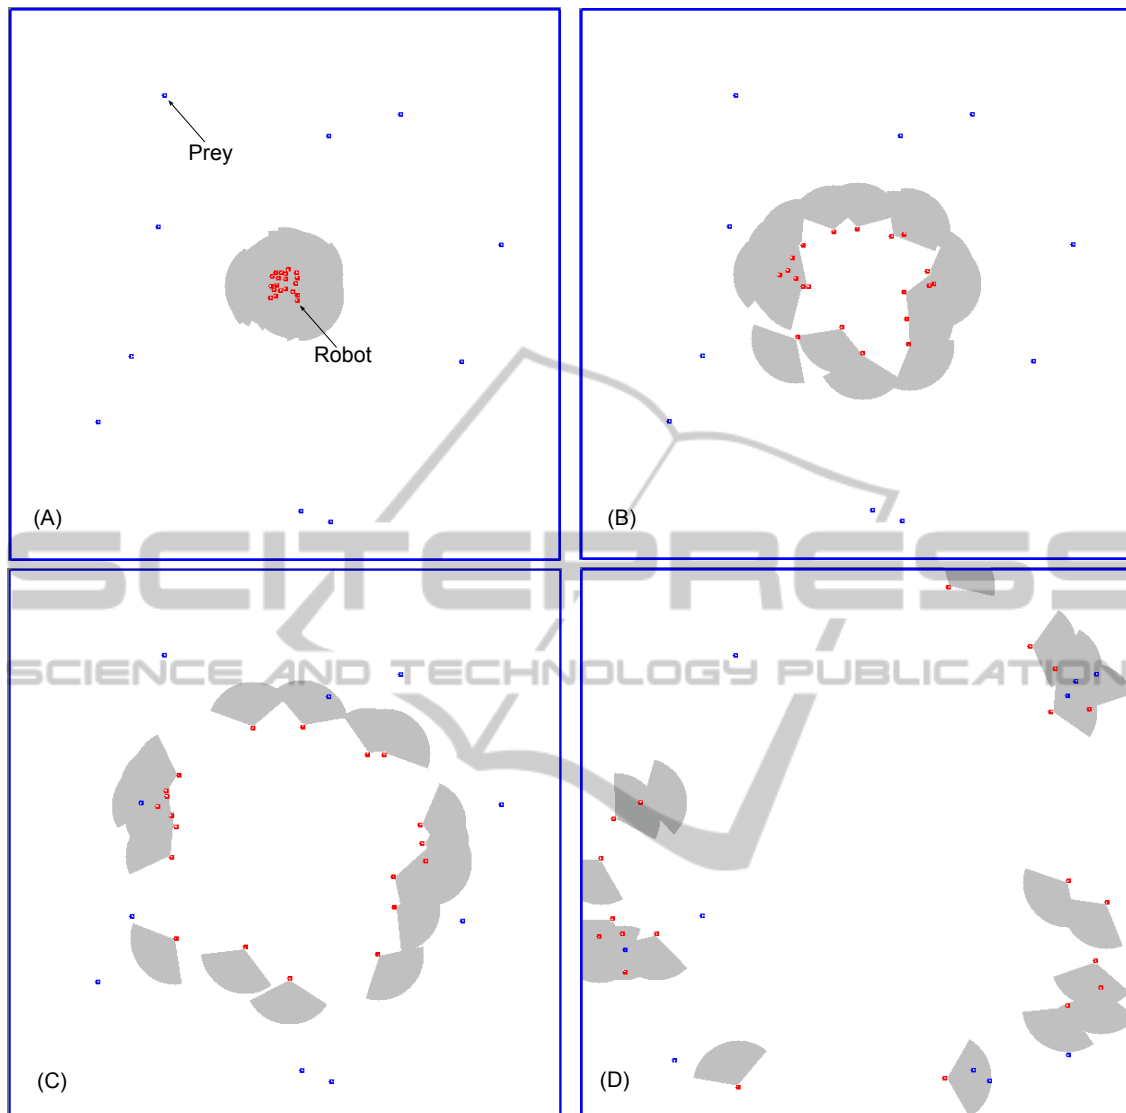


Figure 4: Example of a high performing controller evolved in the collective sensors setup, in a sample with 20 robots and 10 preys. The robots start at the center of the arena (A), and then disperse in outward circular motion in order to find preys (B). When preys are detected (C), the robots chase the preys toward the corners or walls, either in groups or alone (D).

distant from each other, forcing each robot to try to trap a prey alone, or try to maintain the prey in view and wait until another robot gets within range of the collective sensors. In the other five evolutionary runs, the highest-performing controllers of the collective sensors setup display a behavior in which the robots move backwards. Moving backwards represents a poor local optimum in which evolution became stuck in early generations. In this case the robots tend to have a relatively fixed motion pattern that, by chance, can cause preys to be trapped in corners and then consumed.

In the highest-performing behaviors evolved in the onboard sensors setup, the robots start with a simi-

lar behavior to the collective sensors setup, dispersing in different directions to find preys. When a prey is found, the robots attempt to pursue it until another robot be able to detect the same prey with its onboard sensors. The highest-performing controllers of the ideal sensors have a different behavior. Since robots with ideal sensors are almost always capable of seeing a prey, they simply follow and try to consume the closest prey without the need of sometimes extensive periods of searching.

Controllers from the collective sensors setup tend to have a performance closer to the one observed in controllers from the ideal sensors when the number of preys is higher than the number of robots (Figure 5).




Figure 5: The figure shows the number of prey consumed by the highest-performing controllers evolved in the onboard, collective and ideal sensor setups in different combinations of number of preys and robots. Each boxplot represents the average number of preys consumed in 100 post-evaluation samples.

Since the robots in the ideal sensors setup can sense preys at a distance of 3 m, they often tend to follow different preys, which makes it difficult to trap and consume them. The robots in the collective sensors setup, on the other hand, tend to follow fewer preys with more robots, due to their ability to share a prey's location with a limited number of neighboring robots.

6 CONCLUSIONS

In this paper, we explored a novel approach in which robots share readings from their sensors with neighboring robots to overcome the often limited capabilities of the individual robot's onboard sensory hardware. We evaluated our approach in a predator-prey scenario, in which detected preys' estimated positions are communicated to neighboring robots. Robots use received estimates to compute the readings for their collective sensors, thereby effectively allowing robots to sense preys at greater distances, and to more quickly converge on the preys.

Our experimental results showed that swarms using our approach achieve a higher performance than swarms in which the robots have to rely exclusively on their onboard sensors. In certain cases, the performance of swarms using collective sensors even approaches the performance of swarms in which robots are equipped with ideal sensors.

The concept of collective sensors proposed in this paper opens several new avenues of research. Obser-

vations made by different robots can be integrated to allow more precise information to be obtained about the environment. It might be beyond the capability of a single robot to, for instance, estimate the velocity, shape, or size of a particular object, but such estimates could be obtained by combining the sensory readings of multiple robots. Moreover, the sharing of sensory information potentially introduces redundancy in a swarm robotics system. Such redundancy could be used to detect faults, and in case of failure in onboard sensors, a robot could continue to contribute by relying on its collective sensors.

ACKNOWLEDGEMENTS

This work was supported by Fundação para a Ciência e a Tecnologia (FCT) under the grants, SFRH/BD/76438/2011, PEst-OE/EEI/LA0008/2013, and EXPL/EEI-AUT/0329/2013.

REFERENCES

- Bassler, B. L. (1999). How bacteria talk to each other: regulation of gene expression by quorum sensing. *Current Opinion in Microbiology*, 2(6):582–587.
- Beckers, R., Holland, O. E., and Deneubourg, J.-L. (1994). From local actions to global tasks: Stigmergy and collective robotics. In *4th International Workshop on the Synthesis and Simulation of Living Systems (ALIFE IV)*, pages 181–189. MIT Press, Cambridge, MA.

- Beer, R. D. and Gallagher, J. C. (1992). Evolving dynamical neural networks for adaptive behavior. *Adaptive Behavior*, 1(1):91–122.
- Chandrasekaran, S. and Hougen, D. F. (2006). Swarm intelligence for cooperation of bio-nano robots using quorum sensing. In *2006 Bio Micro and Nanosystems Conference (BMN)*, pages 104–104. IEEE Press, Piscataway, NJ.
- Cianci, C., Raemy, X., Pugh, J., and Martinoli, A. (2007). Communication in a swarm of miniature robots: The e-puck as an educational tool for swarm robotics. In *Swarm Robotics*, pages 103–115. Springer, Berlin, Germany.
- Correll, N. and Martinoli, A. (2006). Collective inspection of regular structures using a swarm of miniature robots. In *Experimental Robotics IX*, pages 375–386. Springer, Berlin, Germany.
- Duarte, M., Oliveira, S., and Christensen, A. L. (2011). Towards artificial evolution of complex behavior observed in insect colonies. In *15th Portuguese Conference on Artificial Intelligence (EPIA)*, pages 153–167. Springer, Berlin, Germany.
- Duarte, M., Silva, F., Rodrigues, T., Oliveira, S. M., and Christensen, A. L. (2014). JBotEvolver: A versatile simulation platform for evolutionary robotics. In *14th International Conference on the Synthesis & Simulation of Living Systems (ALIFE)*, pages 210–211. MIT Press, Cambridge, MA.
- Einolghozati, A., Sardari, M., and Fekri, F. (2012). Collective sensing-capacity of bacteria populations. In *2012 IEEE International Symposium on Information Theory (ISIT)*, pages 2959–2963. IEEE Press, Piscataway, NJ.
- Floreano, D., Mitri, S., Magnenat, S., and Keller, L. (2007). Evolutionary Conditions for the Emergence of Communication in Robots. *Current Biology*, 17(6):514–519.
- Fredslund, J. and Matarić, M. J. (2002). A general algorithm for robot formations using local sensing and minimal communication. *IEEE Transactions on Robotics and Automation*, 18(5):837–846.
- Garnier, S., Jost, C., Gautrais, J., Asadpour, M., Caprari, G., Jeanson, R., Grimal, A., and Theraulaz, G. (2008). The embodiment of cockroach aggregation behavior in a group of micro-robots. *Artificial Life*, 14(4):387–408.
- Gutiérrez, A., Campo, A., Dorigo, M., Amor, D., Magdalena, L., and Monasterio-Huelin, F. (2008). An open localization and local communication embodied sensor. *Sensors*, 8(11):7545–7563.
- Mondada, F., Bonani, M., Raemy, X., Pugh, J., Cianci, C., Klaptocz, A., Magnenat, S., Zufferey, J.-C., Floreano, D., and Martinoli, A. (2009). The e-puck, a robot designed for education in engineering. In *9th IEEE International Conference on Autonomous Robot Systems and Competitions (ROBOTICA)*, pages 59–65. IPCB: Instituto Politécnico de Castelo Branco.
- Nolfi, S. and Floreano, D. (2000). *Evolutionary robotics: The biology, intelligence, and technology of self-organizing machines*. MIT Press, Cambridge, MA.
- Rodrigues, T., Duarte, M., Oliveira, S. M., and Christensen, A. L. (2014). What you choose to see is what you get: An experiment with learnt sensory modulation in a robotic foraging task. In *16th European Conference on the Applications of Evolutionary Computation (EvoApplications)*, pages 789–801. Springer, Berlin, Germany.
- Schmickl, T. and Crailsheim, K. (2008). Trophallaxis within a robotic swarm: bio-inspired communication among robots in a swarm. *Autonomous Robots*, 25(1-2):171–188.
- Stamatis, P. N., Zaharakis, I. D., and Kameas, A. D. (2009). A study of bio-inspired communication scheme in swarm robotics. In *7th International Conference on Practical Applications of Agents and Multi-Agent Systems (PAAMS)*, pages 383–391. Springer, Berlin, Germany.
- Støy, K. (2001). Using situated communication in distributed autonomous mobile robotics. In *7th Scandinavian Conference on Artificial Intelligence (SCAI)*, pages 44–52. IOS Press, Amsterdam, The Netherlands.
- Trianni, V., Groß, R., Labella, T. H., Şahin, E., and Dorigo, M. (2003). Evolving aggregation behaviors in a swarm of robots. In *7th European Conference on Artificial Life*, pages 865–874. Springer, Berlin, Germany.
- Turgut, A. E., Çelikkanat, H., Gökçe, F., and Şahin, E. (2008). Self-organized flocking with a mobile robot swarm. In *7th International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, pages 39–46. IFAAMAS, Richland, SC.
- Werfel, J., Petersen, K., and Nagpal, R. (2014). Designing collective behavior in a Termite-Inspired robot construction team. *Science*, 343(6172):754–758.