

Evolving Aggregation Behaviors in a Swarm of Robots

Vito Trianni¹, Roderich Groß¹, Thomas H. Labella¹,
Erol Şahin², and Marco Dorigo¹

¹ IRIDIA – Université Libre de Bruxelles, Belgium

² Dept. of Computer Engineering – Middle East Technical University,
Ankara, Turkey, {vtrianni,rgross,hlabella,mdorigo}@ulb.ac.be
erol@ceng.metu.edu.tr

Abstract. In this paper, we study aggregation in a swarm of simple robots, called *s-bots*, having the capability to self-organize and self-assemble to form a robotic system, called a *swarm-bot*. The aggregation process, observed in many biological systems, is of fundamental importance since it is the prerequisite for other forms of cooperation that involve self-organization and self-assembling. We consider the problem of defining the control system for the *swarm-bot* using artificial evolution. The results obtained in a simulated 3D environment are presented and analyzed. They show that artificial evolution, exploiting the complex interactions among *s-bots* and between *s-bots* and the environment, is able to produce simple but general solutions to the aggregation problem.

1 Introduction

Aggregation is a task of fundamental importance in many biological systems. It is of particular importance for the creation of functional groups of individuals, because it is at the basis of the emergence of various forms of cooperation. In fact, it can be considered a prerequisite for the accomplishment of many collective tasks. These reasons motivate our interest in the study of aggregation within the SWARM-BOTS project, whose aim is the development of a new robotic system, called a *swarm-bot* [10]. A *swarm-bot* is defined as a self-organizing, self assembling artifact composed of a swarm of *s-bots*—mobile robots with the ability to connect to/disconnect from each other. *S-bots* are equipped with simple sensors and motors and have limited computational capabilities. Their physical links are used to self-assemble into a *swarm-bot* able to solve problems that cannot be solved by a single *s-bot*.¹

The SWARM-BOTS project takes inspiration from recent studies in *swarm intelligence*, a novel approach to the design and implementation of “intelligent” systems inspired by the effectiveness and robustness observed in social insects and in other animal societies [3]. A key role in swarm intelligence is played by the

¹ Details regarding the hardware and simulation of the *swarm-bot* are presented in [9] and on the project web-site (<http://www.swarm-bots.org>).

phenomenon of *self-organization*, whereby global level order emerges in a system from the interactions happening among the system's lower-level components. Moreover, these interactions are based only on local information, without reference to the global pattern [5]. A particular form of self-organization is observed in some social insects, which connect one to the other creating complex physical structures. This type of self-organization is referred to as *self-assembling* [1], and is one of the key elements of the SWARM-BOTS project.

Our interest in aggregation stems from the fact that, in order to be able to self-assemble, *s-bots* have to first aggregate in a common location. Additionally, we are interested in self-organized aggregation, that is, aggregation processes that are not driven by a central controller. Self-organization can lead to robust control systems that can be executed without the need of global information, but exploiting only local sensing and local interactions among *s-bots*.

We consider the problem of defining the control system for the *s-bots* using *artificial evolution*, which has gained more and more attention for robotic tasks in the last decade [8]. There are many motivations behind the use of evolutionary techniques for the design of a control system for a robot. In particular, in a multi-robot domain such as the one considered within the SWARM-BOTS project, the dynamical interactions among robots and between robots and the environment make it difficult to hand-design a control system. On the contrary, evolution can fully exploit these dynamic features, without requiring much intervention from the designer. In this paper, we show how evolution can find simple but effective behaviors, which, in some cases, scale with the number of *s-bots* involved in the experiment.

The paper is organized as follows: Section 2 presents some examples of aggregation observed in biological systems, describing the basic mechanisms that enable self-organized aggregation to emerge. Section 3 presents our experiments on evolving aggregation behavior for the *swarm-bot*. The experimental setup is described and the obtained results are analyzed. Finally, Section 4 concludes the paper.

2 Aggregation in Biological Systems

Self-organized aggregation occurs in biological systems by means of two basic mechanisms: *positive* and *negative feedback*. Positive feedback usually takes the form of attraction toward a given signal source (e.g., chemical, tactile, visual). This mechanism leads to amplifications of the source of information, which becomes more and more attractive in time. On the other hand, negative feedback serves as a regulatory mechanism, providing some form of repulsion among the system components, thus controlling the formation of the aggregate.

One of the best studied examples of self-organized aggregation is the one observed in the cellular slime mold *Dictyostelium discoideum* [4,11]. When the amoebae have enough food, they act independently, grow and reproduce rapidly. However, when they are starving, they enter a *developmental phase*: the amoebae emit a chemical attractor that diffuses in concentric waves and serves as a guide

for the aggregation. When the amoebae are clustered, they form a multicellular organism called a slug, which is able to move on the substrate for some time. Eventually, the slug turns into a fruiting body that can develop and distribute spores in the environment, thus restarting the life cycle.

A similar aggregation process can be observed in many other unicellular organisms [5]. Also social and pre-social insects present multiple forms of aggregation, a particular example being the feeding clusters of larvae of the bark beetle *Dendroctonus micans* [6]. In this case, larvae emit pheromone when feeding. The pheromone diffuses in air and triggers the aggregation process. In fact, in presence of a pheromone gradient, larvae react by moving in the direction of higher concentration of pheromone, eventually forming a cluster.

Other interesting examples are given by honey bees, that cluster around the queen on a branch, while scout bees search for new nesting sites. Also birds, fish or mammals present aggregation phenomena that are self-organized, at least to some extent: for example, young penguins aggregate for warmth, many fish species create defensive or hunting schools, and some species of birds and mammals protect their cubs at the center of the group [5].

The aggregation processes described above have been shown to be self-organized. However, we know many other examples in which environmental information or other heterogeneities are used as a clustering stimulus, sometimes mixing it with self-organized behaviors. For example, the environment can offer cues for aggregation, like light or temperature for flies or humidity for sow bugs [5]. However, as mentioned before, we will focus on self-organized aggregation, showing how it can emerge as a result of artificial evolution.

3 Evolving Aggregation Behaviors

As discussed in the previous section, aggregation can be the result of a self-organizing process. After some experience in hand-coding pattern formation behaviors for a group of simulated *s-bots* placed in a grid world [10], we tried to develop *s-bots* able to aggregate by using artificial evolution. In the following, we describe the experimental setup used for the evolution of the clustering behavior. Then, the obtained controllers are analyzed and their properties and limitations are presented.

3.1 Experimental Setup

The experiments presented in this section are performed in simulation, using a software based on the rigid body dynamics simulator SDK VortexTM. This simulator reproduces the dynamics, friction and collisions between physical bodies. The *s-bots* are modeled as cylinders (radius $r = 12$ cm, height $h = 6$ cm) and are provided with two motorized wheels, a gripper that allows connections between *s-bots*, and an omni-directional speaker that continuously produces a tone that can be perceived from a distance up to 50 cm (see Fig. 1a and 1b). Each *s-bot*

is also equipped with eight infrared proximity sensors, three directional microphones, three sensors for detecting established connections on the body, and a gripper sensor simulating a light barrier on the gripper, which is used to perceive the presence of a grippable object (see Fig. 1c). The environment consists of a square arena surrounded by walls. The size of the arena is chosen to be 2×2 meters and is bigger than the perceptual range of the *s-bots* to emphasize the locality of sensing.

We used a generational evolutionary algorithm for the evolution of the *s-bot* neural controller. The genotype specifies the connection weights of a simple perceptron having 17 sensory neurons that encode the state of the 16 sensors and a bias unit (i.e., a unit whose activation state is always 1.0). Each sensory neuron is directly connected to 3 motor neurons, that control the gripper and the speed of the two wheels. The transfer function of the motor neurons is a standard logistic function. Each connection weight ranges in the interval $[-10, +10]$ and is represented in the genotype with 8 bits. Each genotype is mapped into a neural network that is cloned in every *s-bot* involved in the experiment [2]. Five *s-bots* compose the group and they are allowed to “live” for 10 “epochs” (each epoch consists of 600 cycles and each cycle simulates 100 ms of real time). At the beginning of each epoch the *s-bots* are placed in randomly selected positions and orientations within the arena.

In each epoch e , the fitness of a given genotype is estimated averaging over the last 100 cycles a measure $f_e(t)$ that describes the average distance of the group from its center of mass:

$$f_e(t) = \frac{1}{n} \sum_{i=1}^n \left(1 - \frac{d_i(t)}{50} \right),$$

where n is the number of *s-bots* and $d_i(t)$ is the distance of the i^{th} *s-bot* from the center of mass, limited to 50 cm as upper bound in order to have fitness values in the interval $[0, 1]$. The final fitness is obtained by simply averaging the values obtained in each epoch.




Fig. 1. The simplified *s-bot* model. (a) 3D model, where the cylindrical body is transparent to visualize the wheels (b) Actuators: motorized wheels (two gray rectangles), gripper (dotted rectangle), and omni-directional speaker (black circle). (c) Sensors: proximity sensors (black rectangles), directional microphones (white circles), connection sensors (three regions marked with a dashed line around the body), and a light barrier sensor on the gripper (dotted rectangle).




Fig. 2. Snapshots of the formed aggregate. (a) Static clustering behavior. (b) Dynamic clustering behavior.

The population contains 40 genotypes. The best 8 genotypes of each generation are allowed to reproduce, each generating 5 offspring. The per-bit (flip) mutation rate is $2/L$, where L is the length of the genotype. No recombination is used. Parents are not copied to the offspring population. The evolutionary process lasts 100 generations. The experiment was replicated 10 times by starting with different randomly generated initial populations.

In all replications, good solutions are discovered quite early in the evolution, and are then slowly refined afterward, the fitness always reaching values near to the maximum. Figure 2 shows two snapshots taken at the end of the aggregation process.² In the following, we analyze the obtained results, classifying the evolved behaviors in two different classes.

3.2 Behavioral Analysis

By running the evolutionary experiment, we observed the emergence of two types of strategies: a *static* and a *dynamic* clustering behavior. The former creates very compact and stable aggregates in which *s-bots* do not change their relative positions (see Fig. 2a). The latter creates rather loose but moving aggregates that, as we will discuss, allow scalability of the behavior (see Fig. 2b). In the following, we analyze the most representative examples of both classes.

Static Clustering Behavior. As mentioned before, behaviors that fall into the static clustering category create very compact clusters. When far from other individuals, *s-bots* explore the arena moving backward along a circular trajectory having a diameter bigger than the arena side and avoiding walls. When two *s-bots* get close, the attraction to sound sources becomes predominant, and the trajectories change: the two *s-bots* tend to bounce against each other, due to the interplay between attraction and repulsion originating from sound and infrared sensors respectively. In fact, clusters of two *s-bots* are very unstable. However, during the time spent close to each other, the pair can attract other *s-bots*

² See <http://www.swarm-bots.org/index.php?main=3&sub=35&conpage=c1> for some movies of these behaviors.




Fig. 3. Static clustering behavior: fitness evaluations obtained for different group sizes. The average fitness is drawn as a thick line. Boxes represent the interquartile range of the data, while the horizontal bars inside the boxes mark the median values. The whiskers extends to the most extreme data points within 1.5 of the interquartile range from the box. The empty circles mark the outliers.

which can join the cluster, increasing its size and stability. Finally, when the aggregate is formed, *s-bots* continuously monitor the surrounding by performing small moves or turning on the spot, in order to constantly adjust their position with respect to the other *s-bots*.

Note that the performance of the neural controller with respect to the given fitness measure is maximized by this strategy: *s-bots* are in contact when clustered, thus minimizing the distance from the center of mass. Evolution has exploited one important invariant present in the experimental setup: the number of *s-bots*. In fact, given that 5 *s-bots* are present in the environment, only clusters formed by the majority (that is, 3 *s-bots* or more) are stable, while smaller clusters (2 *s-bots*) easily disband. This suggests that when the group size is increased, it will be difficult to obtain a single cluster, but rather multiple smaller clusters will be formed.

In order to confirm this hypothesis, we analyzed the scalability of the clustering behavior with respect to the number of *s-bots* involved. We repeated the evaluation of the fitness 100 times for different group sizes³. The results plotted in Fig. 3 show that, as expected, the behavior does not scale well with the number of robots. In particular, the best results are achieved with group sizes around 5, showing that the evolved behavior is particularly tuned for these situations. Not surprisingly, the average fitness for a group of 2 *s-bots* is very low, as this cluster is unstable. The group of 3 *s-bots* presents an high variance: this suggests that *s-bots* were not able to form a stable cluster within the limited time, due to the lower density of *s-bots* in the arena. For group size bigger than 5 the performance quickly decreases. This can be explained by the fact that, the more

³ Each fitness evaluation is performed over 1,000 cycles, since 600 cycles were not sufficient for the clustering of larger groups.




Fig. 4. Static clustering behavior: snapshots of the sound fields are shown at three instants in time. Each row corresponds to a simulation of clustering with a different group size (5, 10 and 20 *s-bots*, from top to bottom).

the *s-bots*, the higher the number of small clusters. Additionally, a high density of *s-bots* creates a high intensity of sound throughout the arena, causing all the *s-bots* to whirl in place or to join a near neighbor.

The tone continuously emitted by each *s-bot* creates a *sound field* which can be used to give an approximate indication of the sound attraction forces acting on the *s-bots*. Figure 4 plots the change in the sound field over time for groups of 5, 10 and 20 *s-bots*. In the group of five *s-bots*, a single cluster is formed, while multiple clusters appear for larger group sizes. The high intensity of sound inhibits the exploration behavior of the *s-bots* and makes them join the nearest *s-bot* or cluster, as displayed in the last row of Fig. 4.

Dynamic Clustering Behavior. The dynamic clustering behavior creates loose and moving clusters. Also in this case, a circular trajectory is observed when an *s-bot* is far from the walls and from other *s-bots*. When the *s-bots* sense each other, they aggregate and start moving together to explore the environment, in a sort of “flocking” behavior, which is the result of the interplay of attraction to sound and repulsion from too close *s-bots*. This creates moving clusters which can search and merge with other *s-bots* or clusters. When close to each other, *s-bots* continue to move and change their relative positions. In this way, small clusters can change their shape and move across the arena, having the possibility to join other clusters or attract free *s-bots* and increase the size of the aggregate. This feature makes the dynamic clustering behavior robust with respect to the formation of sub-clusters, since formed clusters can continue to explore the arena.

The scalability analysis, conducted in the same way as for the static clustering behavior, confirms the robustness of the evolved behaviors. Figure 5 shows that the performance of the group decreases almost linearly with the group size. This decrease in performance is not due to an imperfect aggregation, but to the fact that the minimum average distance to the center of the cluster grows with the




Fig. 5. Dynamic clustering behavior: fitness evaluations obtained for different group sizes. The average fitness is drawn as a thick line (see also Fig. 3 for explanation of the graph).

number of *s-bots* due to their physical embodiment. Thus, the evolved behavior is well suited for every group size, as sub-clusters, when formed, can continue the aggregation process.

Figure 6 shows the snapshots of sound fields observed for the dynamic clustering behavior. It is worth noting that, unlike the observations made on static clustering behavior, a high intensity of sound in the arena is not problematic for the movement of the *s-bots*. On the contrary, it seems to serve as a communication medium that guides the clustering.




Fig. 6. Dynamic clustering behavior: snapshots of the sound fields are shown at three instants in time. Each row corresponds to a simulation of clustering with a different group size (5, 10 and 20 *s-bots*, from top to bottom).

4 Conclusions

In this paper, we have described the phenomenon of aggregation in biological systems and the evolution of controllers for a group of simulated robots in order to obtain a similar process. In our experiments, two behavioral strategies emerge from evolution. The static clustering behavior results in high fitness values, but it is tuned for a group of size 5. On the contrary, the dynamic clustering behavior obtains lower fitness values as the formed clusters are less compact. However, clusters move through the environment, leading to a scalable behavior.

The reader might have noted that, although *s-bots* can control their grippers, these are not used. In the presented work, we were mainly interested in the study of self-organized aggregation, and not in self-assembling, which requires physical connections. Accordingly, the chosen fitness function does not encourage the establishment of connections. Nevertheless, connections could have appeared because, creating a connection, two *s-bots* minimize their relative distance. What was observed, however, is that connected *s-bots* were, in most cases, unable to move in a coordinated way, making the formation of clusters difficult if possible at all.

In the literature, it is possible to find some interesting works related to the one presented in this paper. Melhuish et al. [7] studied seeded aggregation and collective movement of minimal simulated agents, using sound signals (*chorusing*) to regulate the group size. However, aggregation is not self-organized, because of the presence of an infrared beacon that serves as aggregation site. Yokoi et al. [12], taking inspiration from cellular slime molds, developed amoeba-like robots composed of connected modules. Here, the main difference with the *swarm-bot* is that, even if each module is autonomous in its control, it remains connected to other modules, lacking the full mobility of *s-bots*.

In conclusion, the obtained results show that evolution is able to find simple but effective solutions to the aggregation problem, mainly exploiting some invariants present in the environment and the complex interactions among *s-bots* and between *s-bots* and the environment. Under these conditions, effective behaviors are difficult to hand-design, as they are an “emergent property of the interaction between the robot and the environment” [8]. Furthermore, it is worth noting how the evolved aggregation mechanisms resemble the ones described in Sect. 2. In particular, the attraction to sound sources serves as a positive feedback mechanism: the higher the intensity of sound perceived, the higher the attraction toward the source, which is consequently amplified. On the other hand, the repulsion between *s-bots* constitutes the negative feedback mechanism: it makes clusters of 2 *s-bots* unstable in the static clustering behavior, and results in the movement of the clusters in the dynamic clustering behavior. The latter strategy scales with the number of *s-bots* because it does not strongly rely on environmental invariants, but is merely a result of the dynamic interaction between the *s-bots*, which makes it more similar to the processes observed in nature.

Future work will exploit the presented results to obtain more complex forms of cooperation, like aggregation around preys in order to collectively retrieve them, or aggregation for coordinated motion on rough terrain.

Acknowledgments. Marco Dorigo acknowledges support from the Belgian FNRS, of which he is a Senior Research Associate. The SWARM-BOTS project is funded by the Future and Emerging Technologies programme (IST-FET) of the European Community, under grant IST-2000-31010. The information provided is the sole responsibility of the authors and does not reflect the Community's opinion. The Community is not responsible for any use that might be made of data appearing in this publication.

References

1. C. Anderson, G. Theraulaz, and J.-L. Deneubourg. Self-assemblages in insect societies. *Insectes Sociaux*, 49:99–110, 2002.
2. G. Baldassarre, S. Nolfi, and D. Parisi. Evolving Mobile Robots Able to Display Collective Behaviours. In Hemelrijk C. K. and Bonabeau E., editors, *Proceedings of the International Workshop on Self-organisation and Evolution of Social Behaviour*, pages 11–22, Monte Verità, Ascona, Switzerland, September 8–13, 2002.
3. E. Bonabeau, M. Dorigo, and G. Theraulaz. *Swarm Intelligence: From Natural to Artificial Systems*. Oxford University Press, New York, NY, 1999.
4. J. T. Bonner. *The Cellular Slime Mold*. Princeton University Press, Princeton, NJ, 1967.
5. S. Camazine, J.-L. Deneubourg, N. Franks, J. Sneyd, G. Theraulaz, and E. Bonabeau. *Self-Organization in Biological Systems*. Princeton University Press, Princeton, NJ, 2001.
6. J.-L. Deneubourg, J. C. Gregoire, and E. Le Fort. Kinetics of the larval gregarious behaviour in the bark beetle *Dendroctonus micans*. *Journal of Insect Behavior*, 3:169–182, 1990.
7. C. Melhuish, O. Holland, and S. Hoddell. Convoying: using chorusing to form travelling groups of minimal agents. *Robotics and Autonomous Systems*, 28(2–3):207–216, 1999.
8. S. Nolfi and D. Floreano. *Evolutionary Robotics: The Biology, Intelligence, and Technology of Self-Organizing Machines*. MIT Press/Bradford Books, Cambridge, MA, 2000.
9. G. C. Pettinaro, I. W. Kwee, L. M. Gambardella, F. Mondada, D. Floreano, S. Nolfi, J.-L. Deneubourg, and M. Dorigo. Swarm robotics: A different approach to service robotics. In *Proceedings of the 33rd International Symposium on Robotics*, Stockholm, Sweden, October 7–11, 2002. International Federation of Robotics.
10. E. Şahin, T.H. Labella, V. Trianni, J.-L. Deneubourg, P. Rasse, D. Floreano, L.M. Gambardella, F. Mondada, S. Nolfi, and M. Dorigo. SWARM-BOTS: Pattern formation in a swarm of self-assembling mobile robots. In A. El Kamel, K. Mellouli, and P. Borne, editors, *Proceedings of the IEEE International Conference on Systems, Man and Cybernetics*, Hammamet, Tunisia, October 6–9, 2002. Piscataway, NJ: IEEE Press.
11. C. Van Oss, A.V. Panfilov, P. Hogeweg, F. Siegert, and C.J. Weijer. Spatial pattern formation during aggregation of the slime mould *Dictyostelium discoideum*. *Journal of Theoretical Biology*, 181:203–213, 1996.
12. H. Yokoi, W. Yu, and J. Hakura. Morpho-functional machine: Design of amoeba like robot based on vibrating potential method. *Robotics and Autonomous Systems*, 28(2–3):217–236, 1999.