Behaviour as a Complex Adaptive System: On the role of Self-Organization in the Development of Individual and Collective Behaviour

Stefano Nolfi

Institute of Cognitive Sciences and Technologies
National Research Council (CNR)
Viale Marx, 15, 00137, Rome, Italy
stefano.nolfi@istc.cnr.it

Abstract

In this paper we will explicit the complex system and adaptive nature of behaviour. The complex system nature of behaviour derives from the fact that behaviour and behavioural properties are phenomena that occur at a given time scale and result from several non-linear interactions occurring at a smaller time scale. Interactions occur in time (i.e. consists of a sequence events in which future interactions are constrained by preceding interactions) and might eventually consists of a vector of concurrent interactions. Moreover we argued that behaviour might involve several emergent dynamical processes, hierarchically organized, that affect each others bottom-up and top-down. The adaptive system nature of behaviour derives from the fact that, due to the very indirect relationship between the properties of the interacting elements and the emergent results of the interactions, behavioural systems can hardly be designed while can be effectively synthesized on the basis of a self-organization process (in which properties emerging from interactions can be discovered and retained through an adaptive process based on exploration and selection). These two claims will be demonstrated in two concrete examples involving mobile robots in which non-trivial individual and collective behaviour have been synthesized through an evolutionary technique.

1. Introduction

A new research paradigm, that has been called Embodied Cognitive Science (Varela, Rosch, and Thompson, 1991; Brooks, 1991; Clark, 1997; Pfeifer and Scheier, 1999), has recently challenged the traditional view according to which intelligence is an abstract process that can be studied without taking into consideration the physical aspects of natural systems. In this new paradigm, researchers tend to stress (1) situatedness, i.e., the importance of studying systems that are situated in an environment (Brooks, 1991, Clark, 1997), (2) embodiment, i.e., the importance of study systems that have bodies, receive input from their sensors and produce motor actions as output (Brooks, 1991; Clark, 1997), and (3) emergence, i.e. the importance of viewing behaviour and intelligence as the emergent result of fine-grained interactions[1] between the control system of an agent including its constituents parts, the body structure, and the environment. An important consequence of this view is that the agent and the environment constitutes a single system, i.e. the two aspects are so intimately connected that a

[1] By fine-grained interactions we means interactions that occur at small time scales (e.g. at the time scale of milliseconds).

description of each of them in isolation does not make much sense (Marturana and Varela, 1980, 1987; Beer, 1995).

In section 1 we clarify why behaviour is a complex adaptive system and we discuss how behavioural systems can be developed. After discussing the advantages of self-organizing over design methods, we present two concrete example of effective and robust behavioural system developed through a self-organizing method based on artificial evolution. The first example concerns the development of the control system for an artificial finger that should be able to discriminate objects' shape on the basis of tactile information (section 2). The second example involve the development of the control system of a group of physically assembled robots that should produce coordinated behaviours (section 3). In section 4, we point the hierarchical organization of behaviour. Finally, in section 5, we draw our conclusions.

1.1 Behaviour as a dynamical process resulting from sequences of fine-grained interactions

Behaviour is a dynamical process resulting from the non-linear interactions between an agent (natural or artificial), its body, and the external environment (including the social environment). As we will see, this implies that behavioural systems (such us mobile robots): (1) are extremely difficult to design from the perspective of an external observer, and (2) can be effectively developed through self-organizing methods (e.g. evolutionary methods) that allow to discover and retain useful behavioural properties emerging from the interactions between agents, their bodies, and the environment.

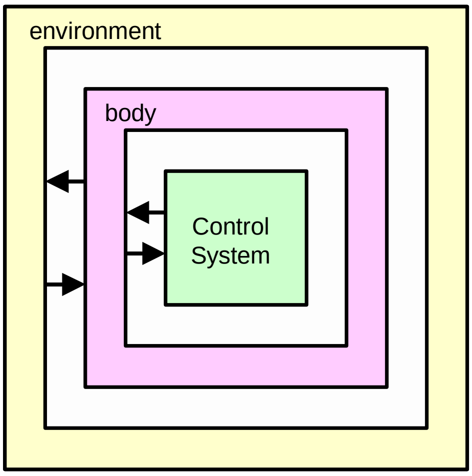

Figure 1. Individual behaviour results from fine-grained non-linear interaction between the agent's control system, its body, and the environment.

At any time step, the environmental structure and the agent/environmental relation influence the body and the motor reaction of the agent that in turn influences the next environmental structure and/or the agent/environmental relation (see Figure 1). Sequences of these form of fine-grained interactions lead to a dynamical process – the behaviour – in which the contributions of the different aspects (i.e. the agent, the body, and the environment) cannot be separated. This implies that even a complete knowledge of the elements governing the interactions provides little insights on the behaviour emerging[2] from the interactions (Maturana & Varela,

[2] We will use the term 'emergence' to indicate a property resulting from a sequence of interactions that can hardly be predicted or inferred from an external observer even on the basis of a complete knowledge of the interacting elements and of the rules governing the interactions.

1980, 1988). The relation between the interaction rules and the resulting behavior is further complicated by the fact that, when interactions are nonlinear, small variations at the levels of the rules governing the interactions might translate to very different forms of behavior due to cumulative and amplifying effects.

1.2 On the advantages of self-organizing over design techniques

From a theoretical point of view, the complex adaptive system nature of behaviour has several important consequences that are far from being fully understood. One important aspect, for instance, is the fact that motor actions partially determine the sensory pattern that agents receive from the environment. By coordinating sensory and motor processes organisms can select favourable sensory patterns and thus enhance their ability to achieve their adaptive goals (Nolfi, 2002, in press; Nolfi & Marocco, 2002; Beer, 2003).

From a engineering point of view, the complex adaptive system nature of behaviour explains why methods based on explicit design are inadequate for developing behavioural systems and why self-organizing methods (e.g. methods based on evolutionary techniques) might be appropriate instead.

The inadequacy of design methods lay on the fact that they require from the designer an ability to infer the rules governing the interactions between the agent and the environment that will lead to a desired behaviour. Unfortunately, as we pointed out above, the properties of the behaviour that emerges from a sequence of fine grained non-linear interactions between the agent and the environment can hardly be inferred from the structure of the interacting elements and the rules governing the interactions. The inverse problem faced by the designer (i.e. the problem of determining the rules governing the interaction that will lead to a desired behaviour) is at least equally hard.

The advantage of self-organizing methods is indeed the fact that they do not require to identify the relation between the rules governing the interactions and the resulting behaviour. They are based on an evolutionary and/or learning process in which the rules governing the interactions, initially randomly assigned, are progressively modified through a process of random variation and selection. Algorithms with this property include evolutionary, simulator annealing, and reinforcement learning algorithms when: (a) the rules governing the interaction are encoded in free parameters, and (b) variations of free parameters are retained or discarded on the basis of variation of performance observed at the behavioural level (i.e. at the time scale of seconds or more). These characteristics allow these methods to discover and retain useful properties emerging from the several interactions without the need to identify the relation between the rules governing the interaction (and/or the interacting elements) and the resulting behaviour.

The possibility to discover and retain useful properties emerging from the interactions also allow self-organizing methods to come up with solutions that are simple from the point of view of the interaction rules (for examples, see Nolfi 2002, in press). Indeed, while in design methods the effects of the detailed characteristics of the agent and the environment (i.e. inertia, elasticity of materials, detailed characteristics of the shape etc.) cannot be predicted and thus constitute problems to be avoided, in self-organizing methods they constitute possibilities to be exploited.

Two example of how self-organizing methods might be used to develop effective behavioural system and to exploit properties emerging from the interactions will be presented in section 2 and 3.

1.3 Collective behaviour emerge from a large number of interactions

Collective behaviour is a dynamical process resulting not only from the fine-grained interactions between agents, their bodies, and the external environment but also between agents (see Figure 2).

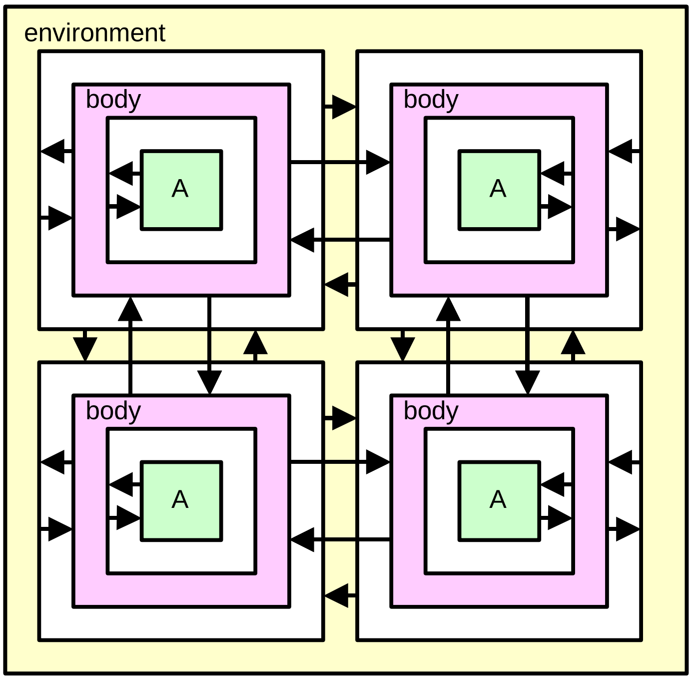

Figure 2. Collective behaviour results from a sequence of several concurrent interactions occurring between agents, their body, and the environment and between agents. The figure schematically represents the situation of four agents in which each individual interacts with two adjacent individuals directly (through physical contact) and indirectly (through environmental modifications that affect other agents' sensors).

The fact that collective behaviour results from a much larger number of fine-grained interactions implies that the relation between the rules governing the interactions and the resulting behaviour is more indirect and more difficult to infer than in the case of individual behaviour. In fact, (a) individual behaviour might be hard to infer or predict on the basis of the rules governing the interactions between the agents, their body, and the external environment (see previous section), (b) groups' aggregate-level behaviour might be hard to infer or predict on the basis of individual behaviours, and (c) the effects of group level dynamics on individual behaviour might be hard to infer or predict. For these reasons, the problem of designing the interaction rules that lead to a desired collective behaviour might be extremely hard even in simple cases (Baldassarre, Nolfi & Parisi, 1993; Funes, Orne & Bonabeau, 2003).

As we mentioned above, however, the indirect relation between the rules governing the interactions and the resulting collective behaviour does not constitute a problem for self-organizing methods. On the contrary the large number of interactions might increase the possibility to identify parsimonious solutions (from the point of view of the complexity of the rules governing the interactions) by exploiting useful behavioural properties emerging from the interactions.

An example of how self-organizing methods might be used to develop effective and robust collective behaviours will be presented in section 3.

2. Evolving the control system of an artificial finger able to discriminate objects with different shapes on the basis of tactile information.

Consider the case of a robot with an artificial finger that has to discriminate objects with different shapes on the basis of rather rough tactile information (Nolfi and Marocco, 2002).

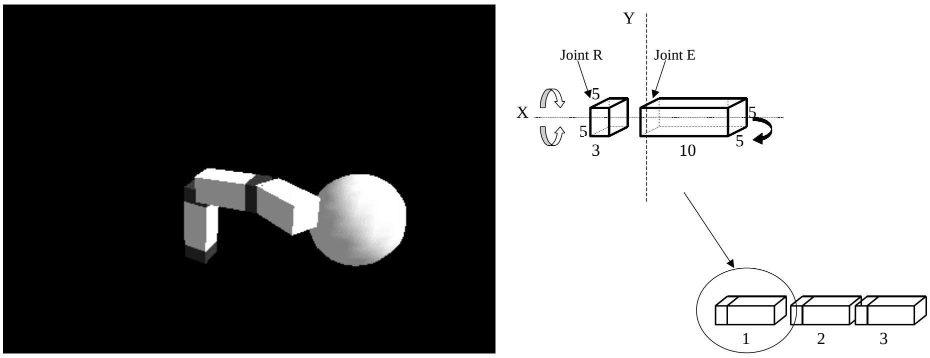

Figure 3. **Left:** The artificial finger and a spherical object. **Right:** A schematic representation of the finger.

The artificial finger consists of 3-segments with 6 degrees of freedom (DOF) and coarse touch sensors (see Figure 3, left). More precisely, the artificial finger consists of a basic structure of two bodies and two joints replicated for three times (see Figure 3, right). These two bodies are connected by means of a joint (i.e. the *Joint E* in Figure 3, right) that allows only one DOF on axis *Y* , while the shorter body is connected at the floor, or at the longer body, by means of a joint (i.e. the *Joint R*) that provides one DOF on axis *X*. In practice, the *Joint E* allows to elevate and to lower the connected segments and the *Joint R* allows to rotate them in both direction. *Joint E and Joint R* are free to moves only in a range between [0 and $\pi/2$] and [$-\pi/2$, $+\pi/2$], respectively. Each actuator is provided with a corresponding motor that can apply a varying force. Therefore, to reach every position in the environment the control system has to appropriately control several joints and to deal with the constraints due to gravity (collisions and physical dynamics was carefully simulated on the basis of Vortex™ libraries).

The sensory system consists of three simple contact sensors placed on each longer body that detect when these bodies collides with obstacles or other bodies and six proprioceptive sensors that provide the current position of each joint. The motor system consists of six motors controlling the corresponding six DOF.

The controller of each individual consists of a neural network with 10 sensory neurons directly connected to 7 motor neurons and 2 internal neurons receiving connections from the sensory neurons and from themselves and projecting connections to the motor neurons. The first 9 sensory neurons encode the angular position (normalized between 0.0 and 1.0) of the 6 DOF of the joints and the state of the three contact sensors located in the three corresponding segments of the finger. The last sensory neuron is a copy of the last motor neuron that encodes the current classification produced by the individual (see below). The first 6 motor neurons control the actuators of the 6 corresponding joints. The output of the neurons is normalized between [0, $+\pi/2$] and [$-\pi/2$, $+\pi/2$] in the case of elevation and rotational joints respectively and is used to encode the desired position of the corresponding joint. The motor is activated so to apply a force proportional to the difference between the current and the desired position of the joint. The seventh motor neuron encodes the categorization output (value below or above 0.5 are interpreted as classifications corresponding to a cubic or spherical object respectively).

The connection weights of the neural controllers were evolved. An initial population of different artificial genotype, each encoding the connection weights of a corresponding neural controller, is created randomly. Each connection weight was represented in the genotype by eight bits that were transformed into a number in the interval [–10, +10]. Each robotic finger is then allowed to interact with the environment on the basis of a corresponding, genetically specified, neural controller. The fittest robots are allowed to reproduce by generating copies of their genotypes with the addition of changes (random mutations). This process is repeated for a number of generations.

Evolving individuals were allowed to "live" for 36 epochs, each epoch consisting of 150 actions. At the beginning of each epoch the finger is fully extended and a spherical or a cubic object is placed in a random selected position in front of the finger (the position of the object is randomly selected between the following intervals: 20.0 >= X <= 30.0; 7.5 >= Y <= 17.5; -10.0 >= Z <= 10.0). The object is a sphere (15 units in diameter) during even epochs and a cube (15 units in side) during odd epochs so that each individual has to discriminate the same number of spherical and cubic objects during its “lifetime” Fitness is computed by counting the number or epochs in which individuals correctly categorize the object (i.e. the number of times in which at the end of the epoch the activation of the last motor units is below 0.5 and the object is a cube or is above 0.5 and the object is a sphere). Therefore, individuals are free to determine how to interact with the objects, i.e. the are only selected on the basis of the ability to correct categorizations.

Population size was 100. The best 20 individuals of each generation were allowed to reproduce by generating 5 copies of their genotype with 1% of their bits replaced with a new randomly selected value.

By running 10 replications of the experiment and by evolving individuals for 50 generations we observed that in many of the replications evolved individuals display a good ability to categorize objects and, in some cases, produce close to optimal performance. Figure 4 shows how a typical evolved individual behave with a spherical and a cubic object (left and right sides of the Figure, respectively). As can be seen, first the finger bends on the left side and move to the right so to start to feel the object with the touch sensor of the third segment. Then the finger continues to move on the same direction by slightly moving up when the third segment of the finger touches the object. As a result of this simple motor rules, in the case of spherical objects, the finger keeps moving toward the left side following the curvilinear surface. In the case of cubic objects, instead, it remains stuck in one of the angles by moving back and fourth.

The behaviour emerging from the interactions between the finger and the objects lead to two rather different behavioural outcomes in the case of spherical and cubic objects: (a) a fully extended position of the finger in the case of spherical objects, and (b) a fully bended position of the finger, in the case of cubic objects. These two positions, in turn, provide a straightforward indication of the type of object the finger interacted with. For other example, involving different environment and robots with different morphologies, in which the convergence or the luck of convergence on a limit cycle behaviour can be used to categorize the environment, see Nolfi, 2002, in press; Nolfi and Marocco, 2002)

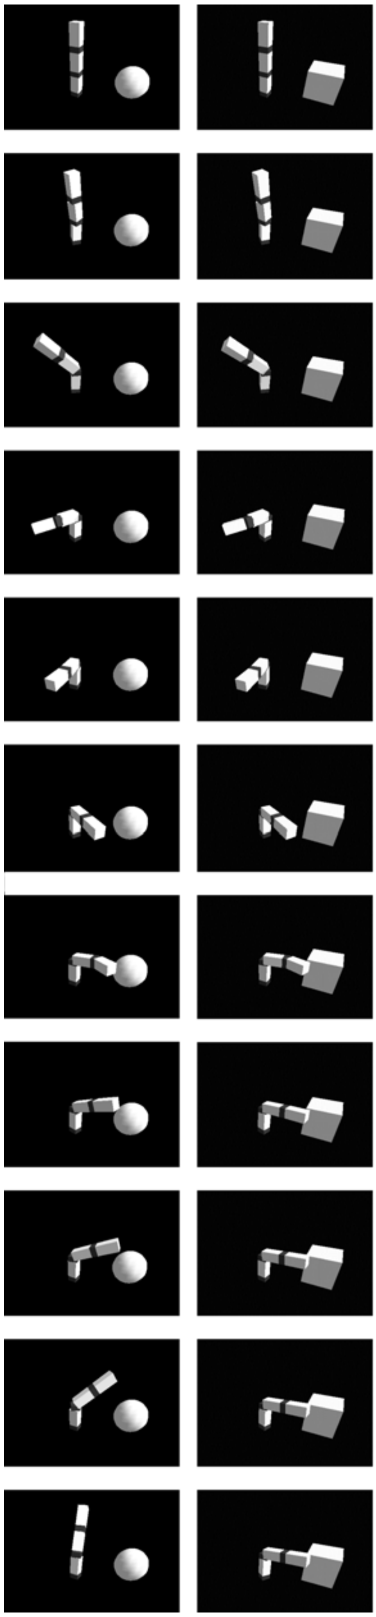

Figure 4. Behaviour of a typical evolved individual during an epoch (150 cycles) in which the object consists of a sphere (left pictures) and of a cube (right pictures). For reason of space, the pictures show the position of the finger each 15 cycles.

Individuals of other replications of the experiments display similar behaviour although the length of the phase with which individuals interact with spherical objects before leaving them varies. The fact that the best performance are observed in cases in which the interaction phase lasts longer (result not shown), demonstrates that the discrimination process is not the result of a single decision but rather the end result of a sequence of interactions between the finger and the object. A similar temporally-extended decision process has been observed by Beer (2003) in evolved agents asked to catch diamonds-shaped objects and avoid circular objects.

3. Evolving the control system of a collection of physically assembled robots able to display coordinated collective behaviour

Consider the case of four assembled robots forming a linear structure (Figure 5) that should move and reach a light target (Baldassarre, Parisi & Nolfi, in press). Given that the orientations of individual robots might vary and given that the target might be out of sight, robots should be able to coordinate to choose a common direction of movement and to change their direction as soon as one or few robots start to detect a light gradient.

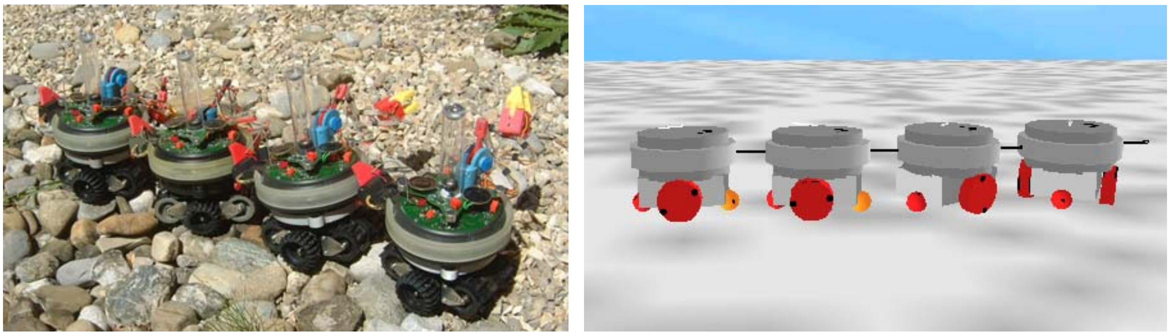

Figure 5. **Left:** Four robots assembled into a linear structure. **Right**: A simplified simulation of the robots described in the left part of the figure based on Vortex^{TM} libraries.

Each robot (Mondada et al., 2004) consists of a mobile base (chassis) and a main body (turret) with a diameter of 116 mm that can rotates with respect to the chassis along the vertical axis. The chassis has two drive mechanisms that control the two corresponding tracks and teethed wheels. The turret has one rigid and one flexible gripper, that allow robots to assemble together and to grasp objects, and a motor controlling the rotation of the turret with respect to the chassis. Robots are provided with a traction sensor, placed at the turret-chassis junction, that detects the intensity and the direction of the force of traction that the turret exerts on the chassis (along the plane orthogonal to the vertical axis) and light sensors. The robots also have several other sensors (a sound sensors, an omnidirectional camera, accelerometers etc.) that, however, were not used in the experiments reported below.

Robots' controller only have access to local sensory information. In particular, each robot's controller consists of a neural network with nine sensory neurons directly connected to two motor neurons. The first four sensory neurons encoded the intensity of the traction from four different orientations with respect to the chassis (rear, left, front and right). The next four sensory neurons provide information on the light gradient with respect to the chassis. The last neuron consists of a bias unit that is always activated to 1.0. The activation state of the two motor neurons was normalized within [–5, +5] rad/s and was used to set the desired speed of the two corresponding wheels and of motor controlling the degree of freedom between the turret and the chassis.

By evolving the connection weights of the robots' controller and by selecting the team of four robots on the basis of the distance travelled from its initial position (when the light target was not on sight) and for the distance travelled toward the target light (when the light target was on sight) we observed that evolving individual are able to effectively solve their problem by negotiating a common direction of movement and by collectively moving toward the light as soon as a light gradient can be detected.

The initial population consisted of 100 randomly generated genotypes that encoded the connection weights of 100 corresponding neural controllers. Each connection weight was represented in the genotype by eight bits that were transformed into a number in the interval [–10, +10]. Each genotype encoded the connection weights of a corresponding neural controllers that was then duplicated four times and embodied into the four robots forming the team (i.e. the team is homogeneous).

By testing evolved controllers in different conditions we surprisingly observed that they are able to generalize their abilities in new conditions and also to spontaneously produce new unexpected behaviours. More precisely, evolved robots display a capacity to generalize their abilities to: (a) the number of assembled robots, (b) the shape with which robots are assembled together, and (c) the use of flexible rather than rigid links. Moreover, evolved robots also display an ability to: (a) spontaneously produce a collective obstacle avoidance behaviour, (b) dynamically rearrange the physical shape of the team in interaction with the environment to negotiate narrow passages, (c) spontaneously produce a coordinate object pushing/pulling behaviour when assembled to or around an external object.

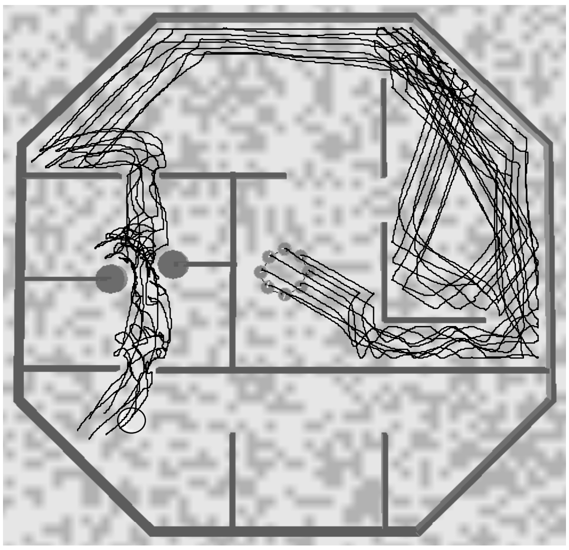

Figure 6. A circular shape structure formed by eight robots assembled through flexible links in a maze with obstacles consisting of walls and cylindrical objects (represented with grey lines and circles). The team of robots starts in the central portion of the maze and reach the light target located in the bottom-left side of the environment (see the light grey circle) by exhibiting a combination of collective obstacle avoidance and collective light approaching behaviour. The irregular lines, that indicate the trajectories of the single robots, provide an indication of how the shape of the assembled robots changes during motion by adapting to the local structure of the environment.

Figure 6, that shows the behaviour displayed by eight robots assembled into a circular shape through flexible links (i.e. links that allow two connected robots to modify their relative positions with limits) placed in a maze environment with walls and cylindrical obstacles, demonstrates how the same control system evolved to control four robots assembled into a linear structure generalize to: (1) a team consisting of eight robots forming a different shape, (2) robots assembled through flexible links that modify the shape of the assembled structure during motion. The figure also show how robots: (a) produce a collective obstacle avoidance behaviour (as a result of the traction force generated during collisions with obstacles), and (b) rearrange the shape of the team to pass narrow passages.

Figure 7, that shows the behaviour of how 8 robots assembled through flexible links around a cylindrical object, demonstrates how the same control system evolved to control four robots assembled into a linear structure generalizes in new conditions and display a co-ordinate object pushing-pulling behaviour.

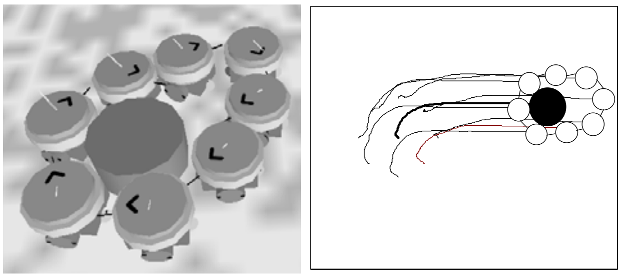

Figure 7: **Left**: Eight robots connected around an object. **Right**: Coordinated object pushing/pulling behaviour exhibited by a team of robots assembled around an external cylindrical object. The empty circles and the full circle indicate the final positions of the robots and of the object, respectively. The thin lines and thick line indicate the trajectory of the robots and of the object.

For a demonstration of how the neural controller evolved in simulation are able to display similar behaviours when embodied and tested in the real physical robots see (Baldassarre, Trianni, Dorigo & Nolfi, in preparation).

4. Behaviour as dynamical system organized hierarchically

In the introduction we pointed out that behaviour is a dynamical process emerging from the interactions between agents' control systems, agents' body, and the external environment (eventually including the social environment). The fact that behaviour (even in simple cases such us grasping an object or reaching a target location) is a property that can be observed only at macro time scale (in the range of seconds or minutes) while interactions occur at micro time scales (milliseconds) imply that behaviour emerge from a large number of non-linear interactions not only in the case of collective behaviour but also in the case of individual behaviour. Behaviour is always the result of a sequence of fine-grained interactions (distributed in time) and eventually of a number of sequential concurrent interactions, occurring between different concurrent sensory-motor processes or between different agents (distributed in space). Overall this imply that individual behaviour and not only collective behaviour is the emergent result of a large number of fine-grained interactions. Although this fact is widely recognized in the case of collective behaviour, it is much less recognized in the case of individual behaviour.

The picture is further complicated by the fact that behaviour might be based on a series of emergent dynamical processes, hierarchically organized, that affect each others bottom-up and top-down (for a similar view, see Keijzer F., 2001). More precisely: (a) interactions between properties emerging from a sequence of fine-grained interactions might lead to higher level emergent properties (that typically extend over larger time scales than the interacting properties), (b) higher level properties might affect the interactions between lower level properties.

As an example of a top-down effects of high level properties (emerging from the interaction between the agent and the environment) and the interaction between the agent and the environment consider the example of the discrimination behaviour described in section 2. The behavioural properties emerging from the interactions between the agent control system, its body, and the external object (occurring at a time scale of 100ms) result in two different emerging behaviours (in the case of cubic or spherical objects respectively): (1) the finger remains bended and keeps touching the object, or (2) the finger becomes fully extended by passing over the object. These two emergent properties occur at a time scale of seconds while the interaction between the agent and the environment are mediated by control rules that operate at the time scale of milliseconds.. These two high level properties, in turns, affect the lower level interactions mediated by the agent neural controller (i.e. the neural controller produces a categorization output corresponding to “cubic object” or “spherical object” on the basis of the state of the sensors that detect the current angular position of the joints of the finger).

As an example of behaviours organized in three hierarchical levels and in which level 3 properties emerge from the interaction between level 2 properties, that in turn emerge from the interaction between the agent and the environment, consider the case of the collective navigation problem described in section 3. Interactions occurring between the agents and with the environment (at a time scale of 100ms) lead to two behavioural properties (that extend at a time scale of seconds): (1) an ability to negotiate and converge on a common direction of movement, and (2) an ability to turn toward the light. The interactions between these two high level properties, in turn, lead to several collective behaviours that occur at larger time scales (i.e. several seconds). More precisely, the interaction between these two behavioural capacity lead to: (a) an ability to collectively approach the light target (even when only few agents detect the light because of their relative distance with respect to the light or because of shadows), (b) an ability to display a collective exploration behaviour and a collective light approaching behaviour and an ability to combine the two behaviours by avoiding to get stuck in situations in which these two behavioural capacity, by triggering opposite motor responses, might interfere one with the other.

5. Conclusion

In this paper we pointed out the complex system and adaptive nature of behaviour.

The complex system nature of behaviour derives from the fact that (both in the case of individual and collective behaviour) behaviour and behavioural properties are phenomena that occur at a given time scale and result from several non-linear interactions occurring at a smaller time scale. Interactions occur in time (i.e. consists of a sequence events in which future interactions are constrained by preceding interactions) and might eventually consists of a vector of concurrent interactions. Moreover we argued that behaviour might involve several emergent dynamical processes, hierarchically organized, that affect each others bottom-up and top-down.

The adaptive system nature of behaviour derives from the fact that, due to the very indirect relationship between the properties of the interacting elements and the emergent results of the interactions, behavioural system can hardly be designed while can be effectively synthesized on the basis of a self-organization process (in which properties emerging from interactions can be discovered and retained through an adaptive process based on exploration and selection).

These two claims have been demonstrated in two examples in which non-trivial individual and collective behaviour have been synthesized through an evolutionary technique.

Acknowledgements

This research has been partially supported by the ECAGENTS project founded by the Future and Emerging Technologies program (IST-FET) of the European Community under EU R&D contract 001940. 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 may be made of data appearing in this publication.

References

Beer R.D. (1995). A dynamical systems perspective on agent-environment interaction. *Artificial Intelligence,* 72:173-215.

Beer, R.D. (2003). The dynamics of active categorical perception in an evolved model agent. *Adaptive Behavior*, 11(4):209-243.

Baldassarre G., Nolfi S., Parisi D. (2003). Evolving mobile robots able to display collective behavior. *Artificial Life*, 9: 255-267.

Baldassarre G., Parisi D., Nolfi S. (in press). Distributed Coordination of Simulated Robots Based on Self-Organisation, *Artificial Life*.

Baldassarre G., Trianni V., Dorigo M. & Nolfi S. (in preparation). Coordinated motion and in teams of physically linked robots.

Brooks R.A. (1991) Intelligence without reason. In J. Mylopoulos & R. Reiter (Eds.), *Proceedings of 12th International Joint Conference on Artificial Intelligence*. San Mateo, CA: Morgan Kaufmann.

Clark A. (1997) *Being There: Putting Brain, Body and World Together Again.* Cambridge, MA: MIT Press.

Funes P., Orme B. & Bonabeau E. (2003). Evolving emergent group behaviors for simple humans agents. In P. Dittrich & J. T. Kim (Eds.) Proceedings of the Seven European Conference on Artificial Life (ECAL 2003).

Keijzer F. (2001). *Representation and Behavior*. Cambridge, MA: MIT Press.

Maturana H.R. & Varela F.J. (1980). *Autopoiesis and cognition: the realization of the living*. Dordrecht: reidel.

Maturana H.R. & Varela F.J. (1988). *The three of knowledge: the biological roots of human understanding*. Boston: New Science Library.

Mondada F., Pettinaro G., Guigrard A., Kwee I., Floreano D., Denebourg J-L, Nolfi S., Gambardella L.M., Dorigo M. (2004). Swarm-bot: A new distributed robotic concept. *Autonomous Robots*. vol. 17 (2-3) : 193-221.

Nolfi S. (2002). Power and limits of reactive agents. *Neurocomputing*, 49:119-145.

Nolfi S. (in press). Categories formation in self-organizing embodied agents. In H. Cohen & C. Lefebvre (Eds), *Handbook of Categorization in Cognitive Science*. Elsevier.

Nolfi S., Marocco D. (2002). Active perception: A sensorimotor account of object categorization. In B. Hallam, D. Floreano, J. Hallam, G. Hayes, J-A. Meyer (eds.) *From Animals to Animats 7: Proceedings of the VII International Conference on Simulation of Adaptive Behavior*. Cambridge, MA: MIT Press, pp. 266-271.

Pfeifer R. and Scheier C. (1999) *Understanding Intelligence*. Cambridge, MA: MIT Press

Varela F.J., Thompson E. and Rosch E. (1991) *The Embodied Mind: Cognitive Science and Human Experience*. Cambridge, MA: MIT Press.