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), Rome, Italy

Key Words

Behavioural systems · Self-organization · Adaptive systems

Abstract

In this paper we discuss 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. consist of sequence events in which future interactions are constrained by preceding interactions) and might eventually consist of a vector of concurrent interactions. Moreover, we argued that behaviour might involve several emergent dynamical processes, hierarchically organized, that af-

Simplexus

For half a century, researchers in artificial intelligence have viewed the human mind as a kind of computer. It may be a strange computer, to be sure, wet and fleshy and built of neurons rather than transistors, but its intelligence rests, they have argued, on the same kind of computational logical as any other computer. From this point of view, the supporting medium of intelligence is nothing but pure logic and mathematics, and should be studied as a 'disembodied' phenomenon in the abstract, without worrying about the messy details of this or that particular system.

Today, of course, this view is on the run. Over the past decade or so, researchers have increasingly come to suspect that our continued failure to create truly intelligent devices stems precisely from the error of this disembodied view. No real world intelligence exists 'in the abstract'; in all its many forms, intelligence always resides within biological brains attached to bodies that interact with environments. Nowadays, this mere 'detail' is increasingly taken as the central key to the puzzle. Thinkers such as Francisco Varela and Rodney Brooks have argued that intelligence does not even reside within an individual, but emerges out of the relationship between an individual and its environment. Intelligence, it seems, may have less to do with the principles of logic and 'top-down' reasoning than with the laws to evolution and self-organization.

In this paper, Stefano Nolfi of the Institute of Cognitive Sciences and Technologies in Rome argues that this fundamental change in perspective has important ramifications for the practical engineering of artificially intelligent devices such as mobile robots. As he points out, the top-down design of artificial agents is often hopeless because of the immense difficulty of foreseeing the long-term consequences of a robot's short-term behavioural rules. Agents always interact with some specific environment, and their actions alter that

fect each other 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 they can be effectively developed through self-organizing methods 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 behaviours have been developed through an evolutionary technique.

Copyright © 2005 S. Karger AG, Basel

1 Introduction

A new research paradigm that has been called embodied cognitive science [1–4] 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 [2, 3], (2) embodiment, i.e., the importance of study systems that have bodies, receive input from their sensors and produce motor actions as output [2, 3], and (3) emergence, i.e. the importance of viewing behaviour and intelligence as the emergent result of fine-grained interactions (i.e. interactions that occur at small time scales) between the control system of an agent including its constituent parts, the body structure, and the environment. An important consequence of this view is that the agent and the environment constitute a

single system, i.e. the two aspects are so intimately connected that a description of each of them in isolation does not make much sense [5–7].

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 examples of an 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 the shape of objects on the basis of tactile information (section 2). The second example involves 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 out 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 as 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 make it possible to discover and retain useful behavioural properties emerging from the interactions between agents, their bodies, 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, which in turn influences the next environmental structure and/

environment; meanwhile, they learn and so alter their own behaviour. The result is a highly unpredictable and uncontrollable co-evolution of agent and environment that makes rational design, at the outset, virtually impossible. But ‘rational’ design is not the only design, and Nolfi demonstrates, in two examples, that engineers can harness self-organization to design intelligent behaviour through a process of evolution and learning.

Taking hints from biology, researchers over two decades or more have explored design through artificial evolutionary processes. In effect, the idea is to build an environment in which a system can discover good designs on its own by gradually changing its operating rules to give better results. The power of such an approach has been demonstrated in areas ranging from statistical physics (algorithms based on simulated annealing) to the design of manufacturing processes (genetic algorithms). The ‘miracle’ of this approach is that it can discover good designs without ever solving the ‘micro-to-macro’ problem or understanding how it is that lower scale rules lead to large-scale outcomes, as Nolfi demonstrates on two practical problems.

First, he shows how an evolutionary approach based on genetic algorithms can design a robotic finger that is able to distinguish geometrical objects – in this case, spheres and cubes. The finger (see fig. 3 of the paper) consists of three segments with various sensors and motors that control and monitor its position. They also sense its possible contact with an object. The sensor ‘neurons’ feed into a system of motor neurons and stimulate them to produce output that makes the finger move. How input gets turned into output depends on the strength of the various wiring connections between sensors and motors, and this network of connections represents the finger’s rudimentary controlling ‘brain’. To design an effective brain, Nolfi has turned to an approach based on a genetic algorithm.

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

or the agent/environmental relation (see fig. 1). Sequences of these forms 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 complete knowledge of the elements governing the interactions provides little insight into the behaviour emerging from the interactions [5, 6]. Please notice that 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.

The relation between the interaction rules and the resulting behaviour is further complicated by the fact that, when interactions are non-linear, small variations at the levels of the rules governing the interactions might translate to very different forms of behaviour 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 [8–11].

From an 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 in 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 need 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

He first created a population of brains or controllers with random connection strength, and had each control the finger through a series of trials during which it tried to distinguish spheres from cubes. In general, random connections will not work very well. But Nolfi let the population of brains ‘reproduce’ to create a new generation and gave the more successful controllers – those doing a better job of discriminating the objects – more offspring. Consequently, good brain designs tended to spread in the population at the expense of bad. Quite quickly, this evolutionary process led to controllers with quite impressive and seemingly intelligent behaviour. In a typical example – illustrated in the author’s figure 4 – the finger distinguishes a cube from a sphere on the basis of the distinct patterns of interaction it has with the two. In the first case, the finger touches an edge of the cube and stops; in the second, it finds no edge and runs itself in a smooth arc over the surface of the sphere. No outside intelligence designed these motions; they evolved quite on their own as a good way to discriminate between the objects the finger finds in its environment. Different behaviours and a different brain would have evolved if the environment had contained pyramids and egg-shaped objects.

Nolfi goes on to demonstrate a similar achievement in a slightly more complex context, with a number of robots that can work together. These robots can link together in various shapes, and each will independently seek to move toward a source of light. However, the robots also monitor their relationships with one another – by sensing, for example, the stress forces associated with pulling or pushing on their neighbours. Again using random rules and an evolutionary approach, the robots quickly learned to coordinate themselves to perform quite delicate tasks such as negotiating a maze, even learning to squeeze the group together to pass through narrow openings.

Fig. 2. Collective behaviour results from a sequence of several concurrent interactions occurring between agents, their body, and the environment and between agents. 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 sensors of other agents) is schematically represented.

Demonstrations of this kind bring home the point that intelligence depends less on logic than on experience. Our own minds bear vestiges of the 99% of human history during which our ancestors lived as hunter-gatherers. Cognitive psychologists believe that our intelligent behaviour, whether in learning how to cooperate with others, or in detecting cheaters, often has less to do with conscious reasoning than with deep instincts that guide our thinking. Nolfi's argument is that we should not expect artificially intelligent devices to be any different. They too need to learn and evolve to do things well. Roughly speaking, engineers should not try to engineer artificially intelligent solutions, but to design interactions that allow systems to evolve and to discover their own solutions.

Mark Buchanan

variation and selection. Algorithms with this property include evolutionary, simulator-annealing, and reinforcement learning algorithms when: (1) the rules governing the interaction are encoded in free parameters and (2) 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 example, see 8, 9]. Indeed, while in design methods the effects of the de-

tailed characteristics of the agent and the environment (i.e. inertia, elasticity of materials, or detailed characteristics of the shape) cannot be predicted and thus constitute problems to be avoided, in self-organizing methods they constitute possibilities to be exploited.

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

1.3 Collective Behaviour Emerges 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 fig. 2).

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, (1) 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 section 1.2), (2) the aggregate-level behaviour of the groups might be hard to infer or predict on the basis of individual behaviours, and (3) 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 [12, 13].

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 [10].

The artificial finger consists of 3 segments with 6 degrees of freedom (DOF) and coarse touch sensors (see fig. 3a). More

Fig. 3. a The artificial finger and a spherical object. **b** A schematic representation of the finger.

precisely, the artificial finger consists of a basic structure of two bodies and two joints replicated three times (see fig. 3b). These two bodies are connected by means of a joint (i.e. the joint E in fig. 3a) that allows only 1 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. Joint R) that provides 1 DOF on axis X. In practice, Joint E makes it possible to elevate and to lower the connected segments and Joint R allows to rotate them in both direction. Joint E and Joint R are free to move only in a range between $(0 \text{ and } \pi/2)$ 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 collide 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 6 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 as 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 (values 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 a 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

Fig. 4. Behaviour of a typical evolved individual during an epoch (150 cycles) in which the object consists of a sphere **(a)** and of a cube **(b)**. For reasons of space, the pictures show the position of the finger every 15 cycles.

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 randomly selected position in front of the finger (the position of the object is randomly selected between the following intervals: $20.0 \leq X \leq 30.0$; $7.5 \leq Y \leq 17.5$; $-10.0 \leq Z \leq 10.0$). 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. they 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 behaves with a spherical and a cubic object (fig. 4a and b, respectively). As can be seen, first the finger bends on the left side and moves to the right so as to start to feel the object with the touch sensor of the third segment. Then the finger continues to move in the same direction by slightly moving up when the third segment of the finger touches the object. As a result of these 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 leads to two rather different behavioural outcomes in the case of spherical and cubic objects: (1) a fully extended position of the finger in the case of spherical objects and (2) a fully bent 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 examples, involving a different environment and robots with different morphologies, in which the convergence or the lack of convergence on a limit cycle behaviour can be used to categorize the environment, see 8–10].

Individuals of other replications of the experiments display a 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 is observed in cases in which the interaction phase lasts longer (results not shown) demonstrates

Fig. 5. **a** Four robots assembled in a linear structure. **b** A simplified simulation of the robots described in the left part of the figure based on Vortexlibraries.

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 in a different experimental setup in which evolving agents are asked to catch and avoid objects with different shapes [11].

3 Evolving the Control System of a Collection of Physically Assembled Robots Able to Display Coordinated Collective Behaviour

Consider the case of 4 assembled robots forming a linear structure (fig. 5) that should move and reach a light target [14]. 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.

Each robot [15] consists of a mobile base (chassis) and a main body (turret) with a diameter of 116 mm that can rotate with respect to the chassis along the vertical axis. The chassis has two drive mechanisms that control the two corresponding tracks and toothed wheels. The turret has one rigid and one flexible gripper that allow robots to assemble and to grasp objects, and a motor controlling the rotation of the turret with respect to the chassis. Ro-

bots 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 sensor, an omnidirectional camera, accelerometers) that, however, were not used in the experiments reported below.

The controllers of the robots only have access to local sensory information. In particular, each controller of a robot consists of a neural network with 9 sensory neurons directly connected to 2 motor neurons. The first 4 sensory neurons encoded the intensity of the traction from four different orientations with respect to the chassis (rear, left, front and right). The next 4 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 2 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 DOF between the turret and the chassis. 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 connec-

tion weights of a corresponding neural controllers that was then duplicated 4 times and embodied into the 4 robots forming the team (i.e. the team is homogeneous).

By evolving the connection weights of the controller of the robots and by selecting the team of 4 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 individuals 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.

By testing evolved controllers in different conditions we surprisingly observed

Fig. 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 reaches 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 represent the trajectories of the individual robots provide an indication of how the shape of the assembled robots changes during motion by adapting to the local structure of the environment.

Fig. 7. a Eight robots connected around an object. **b** 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.

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: (1) the number of assembled robots, (2) the shape with which robots are assembled, and (3) the use of flexible rather than rigid links. Moreover, evolved robots also display an ability to: (1) spontaneously produce a collective obstacle avoidance behaviour, (2) dynamically rearrange the physical shape of the team in interaction with the environment to negotiate narrow passages, and (3) spontaneously produce a coordinate object pushing/pulling behaviour when assembled to or around an external object.

Figure 6 shows the behaviour displayed by 8 robots assembled in a circular shape through flexible links (i.e. links that allow 2 connected robots to modify their relative positions within certain limits) placed in a maze environment with walls and cylindrical obstacles. As shown the same control system evolved to control 4 robots assembled into a linear structure generalize to: (1) a team consisting of 8 robots assembled into a different shape and (2) robots assembled through flexible links that modify the shape of the assembled structure during motion. Figure 6 also shows how robots: (1) produce a collective obstacle

avoidance behaviour (as a result of the traction force generated during collisions with obstacles) and (2) rearrange the shape of the team to pass narrow passages.

Figure 7 shows the behaviour of 8 robots assembled through flexible links around a cylindrical object. As shown, the same control system evolved to control 4 robots assembled into a linear structure generalizes in new conditions and displays a coordinate object pushing-pulling behaviour.

For a demonstration of how the neural controller evolved in simulation is able to display similar behaviours when embodied and tested in the real physical robots, see Baldassarre et al. [16].

4 Behaviour as Dynamical System Organized Hierarchically

In the introduction we pointed out that behaviour is a dynamical process emerging from the interactions between the agents' control systems, the agents' body, and the external environment (eventually including the social environment). The fact that behaviour (even in simple cases such as grasping an object or reaching a target location) is a property that can be observed only at a macro time scale (in the range of seconds or minutes) while interactions occur at micro time scales (milli-

seconds) implies that behaviour emerges 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 concurrent interactions between different agents (distributed in space). Overall this implies that both individual behaviour and collective behaviour are 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 series of emergent dynamical processes, hierarchically organized, that affect each other bottom-up and top-down [for a similar view, see 17]. More precisely: (1) 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) and (2) higher-level properties might affect the interactions between lower-level properties.

As an example of top-down effects of high-level properties (i.e. behavioural properties emerging from the interaction between the agent and the environment) on lower levels of agent/environmental interactions 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 100 ms) result in two different emerging behaviours (in the case of cubic or spherical objects, respectively): (1) the finger remains bent 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 is mediated by control rules that operate at the time scale of milliseconds. These two high-level properties, in turn, affect the successive 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, which 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 the environment (at a time scale of 100 ms) 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 capacities leads to: (1) 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), (2) an ability to display a collective exploration behaviour and a collective light-approaching behaviour, and (3) an ability to combine the two behaviours by avoiding to get stuck in situations in which these two behavioural capacities, by triggering opposite motor responses, might interfere with one another.

5 Conclusion

In this paper we focused on 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. consist of sequence events in which future interactions are constrained by preceding interactions) and might eventually consist of a vector of concurrent interactions. Moreover, we argued that behaviour might involve several emergent dynamical processes, hierarchically organized, that affect each other 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, the behavioural system can hardly be designed while it 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 has been developed 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

- 1 Varela FJ, Thompson E, Rosch E: *The Embodied Mind: Cognitive Science and Human Experience*. Cambridge, MIT Press, 1991.
- 2 Brooks RA: *Intelligence without reason*; in Mylopoulos J, Reiter R (eds): *Proceedings of 12th International Joint Conference on Artificial Intelligence*. San Mateo, Morgan Kaufmann, 1991.
- 3 Clark A: *Being There: Putting Brain, Body and World together Again*. Cambridge, MIT Press, 1997.
- 4 Pfeifer R, Scheier C: *Understanding Intelligence*. Cambridge, MIT Press, 1999.
- 5 Maturana HR, Varela FJ: *Autopoiesis and Cognition: the Realization of the Living*. Dordrecht, Reidel, 1980.
- 6 Maturana HR, Varela FJ: *The Three of Knowledge: the Biological Roots of Human Understanding*. Boston, New Science Library, 1988.
- 7 Beer RD: A dynamical systems perspective on agent-environment interaction. *Artif Intell* 1995; 72: 173–215.
- 8 Nolfi S: Power and limits of reactive agents. *Neurocomputing* 2002; 49: 119–145.
- 9 Nolfi S: Categories formation in self-organizing embodied agents; in Cohen H, Lefebvre C (eds): *Handbook of Categorization in Cognitive Science*. Amsterdam, Elsevier, 2005.
- 10 Nolfi S, Marocco D: Active perception: a sensorimotor account of object categorization; in Hallam B, Floreano D, Hallam J, Hayes G, Meyer J-A (eds): *From Animals to Animats. 7. Proceedings of the 7th International Conference on Simulation of Adaptive Behavior*. Cambridge, MIT Press, 2002, pp 266–271.
- 11 Beer RD: The dynamics of active categorical perception in an evolved model agent. *Adapt Behav* 2003; 11: 209–243.
- 12 Baldassarre G, Nolfi S, Parisi D: Evolving mobile robots able to display collective behavior. *Artif Life* 2003; 9: 255–267.
- 13 Funes P, Orme B, Bonabeau E: Evolving emergent group behaviors for simple humans agents; in Dittrich P, Kim JT (eds): *Proceedings of the 7th European Conference on Artificial Life (ECAL 2003)*, Dortmund, 2003.
- 14 Baldassarre G, Parisi D, Nolfi S: Distributed coordination of simulated robots based on self-organisation. *Artif Life*, in press.
- 15 Mondada F, Pettinaro G, Guirard A, Kwee I, Floreano D, Denebourg J-L, Nolfi S, Gambardella LM, Dorigo M: *Swarm-bot: a new distributed robotic concept*. *Autonomous Robots* 2004; 17: 193–221.
- 16 Baldassarre G, Trianni V, Bonani M, Mondada F, Dorigo M, Nolfi S: *Coordinated motion and in teams of physically linked robots*. Technical Report. Rome, Institute of Cognitive Science and Technology (CNR), 2005.
- 17 Keijzer F: *Representation and Behavior*. Cambridge, MIT Press, 2001.