

Robotics in the Science of Complex Systems

Jeffrey Johnson

The Open University, Milton Keynes MK63DA, UK,
j.h.johnson@open.ac.uk,
<http://design.open.ac.uk/johnson>

Abstract. Multi-robot systems may provide an important step towards a general science of complex systems. When classified according to experimental replicability, teams of interacting robots form an intermediate class of systems, more complex than many systems in traditional physical science but less complex than most human systems. It is difficult for humans to study human systems from the inside; and often impossible to perform repeatable experiments on complex human systems. Team robotics does not have these problems, while still exhibiting many of the characteristics of complex systems. It thus provides an ideal laboratory to investigate some of the most fundamental questions in the evolving science of complex systems.

1 Introduction

It will be argued that multi-robot systems and team robotics can be an important step towards a general science of complex systems. Although there is not yet a definitive list, some fundamental questions can be identified as major hurdles to progress in the science of complex systems. Unless they are overcome, the science cannot progress. These fundamental bottleneck problems include

- how can complex systems and their dynamics be represented in computers?
 - how can multi-level lattice-hierarchies be represented in computers?
 - what multi-level management and control structures are in complex systems?
 - how can multi-level self-organisation be represented?
 - how can computationally irreducible dynamics be represented?
 - how can co-evolution dynamics between subsystems be represented?
 - what are the implications of the ontological assumption for complex systems?
- how can computers form their own representations?
 - how can machines abstract their own constructs?
 - how can communities of machines agree vocabulary and language?
 - how can language be grounded in observation?
 - how can pattern recognition be automated?
- what are the scientific foundations of computer simulation?
 - how can we answer the “can you trust it?” problem of computer simulation?
 - how can we answer the “invented universe” problem of computer simulation?
 - what are the limits of massively parallel computation?
 - does computation establish natural limits to prediction and what is knowable?

- how can synthetic complex systems be designed, controlled, and managed?
 - what is the nature of designing and managing synthetic systems?
 - how can we engineer emergence in complex synthetic systems?
 - where are the boundaries of complex systems?
 - how can large heterogeneous inconsistent and incoherent data sets be used?
 - can synthetic systems self-regulate, and form autonomous living systems?
 - what kinds of cultures emerge from machines forming their own languages?

and there are others. Trying to answer questions like these in the context of human systems is particularly difficult for various reasons, including:

- it is difficult for humans to study human systems objectively from the inside;
- it is difficult to control human systems in experiments - humans don't obey orders
- it is difficult to know if subjects are stakeholders in the outcome of experiments
- it is difficult to perform repeatable experiments on complex human systems.

The traditional scientific approach involves observing systems and abstracting their properties. The early stages may involve unsystematic and open-ended experiments to investigate what happens when the system is poked this way or that, and their outcomes may be completely unpredicted. From this follows a description of the system in a symbolic language, usually vernacular language augmented by system-specific terms, mathematics, diagrams, or images. The symbolic language evolves during the observation process. Then observations are synthesised into a theory. Within the theory there may be entailments of relationships not yet observed, leading to hypotheses and experiments to test them and the whole theory. However, this approach may be impractical:

- | | |
|---------------------------|---|
| • logical exclusion | - A wins an election, the result of B winning is unknowable |
| • moral exclusion | - cannot perform harmful experiments on people |
| • risk exclusion | - cannot undertake dangerous experiments |
| • political exclusion | - performing an experiment would be fatally unpopular |
| • financial exclusion | - particular social experiments may be too expensive |
| • temporal exclusion | - experiments take too long (planning a town) |
| • spatial exclusion | - one configuration of building precludes another |
| • sampling exclusion | - impossible to collect enough samples to build distributions |
| • extinction exclusion | - experiments causing extinction pre-empt further experiments |
| • contamination exclusion | - social experiments may change people |
| • unique state exclusion | - the initial conditions may never be repeated |

These constraints need not apply to teams of soccer-playing robots, even though they *are* complex systems with many of the characteristics of complexity, and provide well-defined examples of the fundamental bottleneck problems.

Given that the fundamental questions can be investigated in robot systems, laboratory experiments can be performed without humans being in the loop. In other words, robot systems can be observed from the outside, allowing a clear distinction to be made between the values, culture and aspirations of the robot agents inside the system, and the values, culture and aspirations of the human observers outside the system.

Finally, given the possibility of investigating the fundamental questions without the complication of humans both inside and outside the system, it is argued that robots overcome the exclusions described above. It is possible to set up systematic and repeatable experiments with robot systems that would be impractical or impossible for human systems. The ability to plan experiments with robot systems promises to provide significant insights into the fundamental questions in the relatively short term. Robot teams form an intermediate class of complex systems, less complex than human systems, but subject to many of the fundamental questions of complexity. Thus, robotics can provide a unique laboratory for research into the fundamental problems in the science of complex systems.

2 Robot systems can be studied without the complications of human agents

One of the astonishing characteristics of humans is the ability to formulate predictions on the basis of no explicit prior theory or data: we can encounter new environments, formulate implicit ad-hoc models, and make relatively good decisions on how to act. Often this involves making analogies between diverse domains, and requires a great deal of knowledge in many areas. Eliciting this information can be a complicating factor in trying to understand human systems.

Researchers and analysts of many human systems are often ‘stakeholders’ in the outcome of the research, which can lead to a lack of scientific objectivity. Also humans can be very opinionated, and are not always the best instruments for collecting data. By comparison, robot systems can be studied in very objective ways. Data comes from objective sensors, and can be recorded. The theories on which soccer-playing robots operate are in very limited domains and are much simpler than the kind of theories that humans formulate.

Human beings have wonderful pattern recognition abilities that far surpass anything that can be done by machines, but this can be hard to investigate. By comparison, in robot systems all pattern recognition is automatic, making them easier to interrogate and understand.

Another complicating factor in studying human systems is that social structures act as a backcloth for information flows, including ‘irrelevant normal life’ information. In robot systems information flows need not be degraded by the background clutter of living a normal human life, and depend on explicit connectivities, making their analysis much easier.

Thus, robot systems avoid the complications of human systems because:

- autonomous mobile team robots don’t need to be stakeholders
 - robots don’t need to have egos or worries
 - robots don’t need to feel pride or shame
 - robots don’t need to have to hold down a job
 - robots don’t need to tell lies
 - robots don’t need to cheat
 - robots don’t need to disobey commands in experiments

- autonomous mobile team robots just do what they do
 - and we can observe them
 - all robot observations are made by sensors
 - all data is explicit within the representation
 - all pattern recognition is automatic
 - all information is in an explicit vocabulary
 - robots don't live a life outside experiments
 - humans can observe construct and language formation

3 Societies of robots are complex systems and can be used for repeatable systematic laboratory experiments

While it is not possible to do systematic experiments in many human systems, you can have it both ways in robotics and run the system with and without the change, or across a wide range of possibilities. Also, with robots it is possible conduct systematic series of experiments to investigate evolution and co-evolution between robots and their environment.

Robot systems are less subject to initial state replicability problems than social systems, and their sensitivity to initial conditions can be investigated by careful laboratory experiments. Robots experiments can be run many times to obtain distributions of outcomes - which is often impractical in human systems.

The RoboCup Challenge that a team of humanoid robots will beat the human world champion soccer players by 2050 [13] has stimulated a huge programme of research across many fields, and has central importance in complexity science [7]. The Simulation League is particularly interesting from the perspective of complexity science, being open, robust and well-documented. With teams freed from the engineering problems of physical robots, progress in tactics and strategy has been greatest in this league. The best teams manage to get the simulated robots to work together as multi-agent systems.

Robot soccer players are autonomous agents, and multi-agent teams of robot soccer players are complex adaptive systems, evolving their own self-organising behaviour, and co-evolving through playing with other robot teams. Competing strategies in robot soccer resonate with game theory, increasing returns and path dependence [1]. Robot soccer teams exhibit emergence and chaos with computationally irreducible dynamics. They require explicit symbolic and analogue representations and languages. As such, robot soccer is an excellent laboratory for experiments in complexity science.

4 Team Robotics and the Bottleneck Problems of Complexity Science

4.1 Robots and explicit representation inside computers:

One of the fundamental questions in complexity science is how can we represent complex systems and their dynamics inside computers? Robot soccer is computationally irreducible, and everything must be inside the machine. This requires

some explicit way of representing its data and information, and some explicit way of processing those data.

Robotics and Multilevel Representation:

Human beings perceive and represent the world in terms of hierarchical structures, e.g. our languages contain collective nouns reflecting a set-theoretic hierarchy of inclusion. It is increasingly being recognised that classes may intersect, and that most set hierarchies are non-partitional, *e.g.* water polo belongs to both the set of water sports and the set of ball games [5], *e.g.* a soccer pitch may be divided into intersecting areas, such as those 'attacked' by individuals at any instant of time - opposing players can interact only where the areas intersect.

There are also hierarchies of assembly in which sets of parts are assembled to form new wholes [6]. These are of particular significance in complexity science, since sets are assembled to form wholes with *emergent properties*, *e.g.* components are assembled to form aeroplanes, which then possess emergent properties such as being able to fly. When artifacts are designed, component parts are specified, as are the assembly relations that say how the parts must be put together. Sometimes the emergent properties of the whole can be deduced from the properties of the components, as in electrical engineering design, and sometimes they are a complete surprise, as with Chernobyl.

In robot soccer, players assemble themselves to form attacking structures, or defending structures. These can be highly dynamic, with the structures forming and breaking up as the players move. It must be possible for these structures to be represented inside computers, so that robots can recognise and reason about them [8].

Assembling a specified set of components in a specified way may also connect those components with other unspecified elements of a system, producing unexpected outcomes. A major challenge in complexity science is understanding how connectivities between subsystems through lower level substructures induce unexpected interactions [6].

Hierarchies of authority and control:

There is a difference between the hierarchies discussed in the previous section, and hierarchies of authority in human systems where one person can tell another person what to do. A distinction is then made between top-down hierarchies, in which bosses at every level tell lesser bosses what to do, and bottom-up hierarchies in which collections or teams of agents self-organise autonomously to decide what to do.

In human organisations, the mission or goals of the organisation is decided 'at the top'. Without this top-down control, the system would lack coherence, and nothing could be achieved. Of course there is self-organisation at every level, and local goals may be questioned at every level, and they may be modified as projects proceed. The issues of goal and subgoal determination can be investigated using teams of robots [8], and in this respect, multi-robot systems provide an excellent laboratory for investigating hierarchies of authority.

Computationally irreducible dynamics and co-evolution

Robot systems are very sensitive to initial conditions. Robot system dynamics are chaotic and computationally irreducible, and the success or failure of alternative control approaches can be clearly observed in the laboratory - especially at higher representational levels.

In robot soccer, the tactics and strategies used by teams change according to the performance of other teams. If team A beats team B consistently, then team B must change the way it plays. If this results in team A losing consistently, it too will have to change the way it plays. This process of co-evolution can go on indefinitely. This provides a relatively simple laboratory subject for the study of co-evolutionary dynamics in complex systems.

Machines that abstract their own representations

In general complex systems involve the collection great quantities of heterogeneous data. Traditionally, humans build the lattice hierarchical vocabulary, and map data onto this representation. This raises two questions: do humans produce the most appropriate vocabulary, and do humans need to do all the pattern recognition?

For example, in machine vision, it seems highly likely that humans will not produce the most appropriate structures for recognising objects, and that it might be better to let machine vision systems learn their own primitive features and intermediate structures. It is widely accepted that biological vision involves a series of levels[14]. These are related to the lattice hierarchies discussed previously. A single pixel conveys little information, while, structured sets of pixels convey more. Although terms like ‘edge’ seem very natural, such constructs may not be the best way to represent objects in machine vision, and it is interesting to ask if machine vision systems can abstract their own constructs, taking humans out of the loop [11].

Automatic pattern recognition is one of the most fundamental bottleneck problems in the science of complex systems. In human systems many pattern recognition operations are done by people, and we don’t know how they do it.

There are many problems in robotics that can also be used to investigate automatic construct abstraction for pattern recognition. These include robot soccer player recognising good positions for scoring goals, and recognising good trajectories to attain those positions [8, 10]. Thus, robotics provides an excellent laboratory, with repeatable experiments, for the study of automated pattern recognition.

The ontological assumption and complex systems:

The representation problem often involves finding appropriate ontologies, *i.e.* ways of representing classes of things, their properties, and relationships between things. It seems that symbolic representation is crucial in ‘putting systems inside machines’, but Dreyfus writes that “after fifty years of effort, however, it is now clear to all but a few diehards that this attempt to produce general intelligence has failed. This failure does not mean that this sort of AI is impossible; no one

has been able to come up with such a negative proof. Rather, it has turned out that, for the time being at least, the research program based on the assumption that human beings produce intelligence using facts and rules has reached a dead end, and there is no reason to think it could ever succeed.” [4]. However, while information processing in neural-like systems is very important in creating ‘intelligent’ machines, it needs to be set alongside symbolic representation and computation, analogous to right-left brain complementarity

Although complex systems scientists are not necessarily interested in building human-like intelligent systems, the necessity to represent complex systems inside computers creates many parallels. Dreyfus’ says that there is no evidence that the mind processes information in a sequence of discrete steps like a heuristically programmed digital computer, or that human behaviour can be formalized in these terms. This includes four types of human information processing which resisted formalisation in terms of heuristic rules: fringe consciousness, ambiguity tolerance, essential/inessential discrimination, and perspicuous grouping. Dreyfus also suggests that data cannot be available in the necessary form: “the data with which the computer must operate if it is to perceive, speak, and in general behave intelligently, must be discrete, explicit, and determinate; otherwise it will not be the sort of information which can be given to the computer so as to be processed by rules. Yet there is no reason to suppose that such data about the human world are available to the computer and several reasons to suggest that no such data exist.” “The ontological assumption that everything essential to intelligent behaviour must in principle be understandable in terms of a set of determinate independent elements allows AI researchers to over look this problem” [4, page 206].

Robot systems give a powerful way of investigating these ideas. Robots do use rules and symbolic reasoning in their representation, but they may also use neural processing. Furthermore, robots have embodied intelligence, in which part of their behaviour is based on analogue processing in the materials of which they made. And for experiments into the fundamental questions, it may not matter whether robot ‘intelligence’ is human-like.

4.2 Robot simulation and complexity science

Simulation is a particularly important feature of robot soccer. In robotics it is common to test programs by graphical simulators, before new strategies are programmed in to the physical robots. At its simplest, simulation allows programs to be debugged, saving time diagnosing and fixing programming errors in hardware. At higher levels, robot soccer is characterised by discrete dynamics, with cellular automata being widely used to simulate the outcome of tactics and strategies in the lattice hierarchy of representation.

Robot simulation and the ‘Can you trust it?’ problem:

Simulations are used extensively in complex systems science, but they are subject to what has been called ‘the can-you-trust-it?’ problem [3, 9]. Put simply,

what are the criteria for accepting that a simulation will give observable results? Currently there is no definitive test to know if a simulation is good or not.

Simulations of complex systems usually involve two parts: a simulation of the system, and a simulation of its environment. The ‘can you trust it?’ problem therefore has two components:

- can you trust the simulation of the system?
- can you trust the simulation of the system’s environment?

There is further room for doubt:

- can you trust the data?
- can you trust the processing?

Simulating and calibrating the environment can be as difficult as simulating and calibrating the system. It is easiest in physical systems, which is why simulations based on Finite Element Analysis *are* widely trusted by engineers. The environment for mechanical and electrical systems is relatively well understood, thanks to five hundred years research in physical science.

By comparison, the environment for many social systems is poorly understood. What constitutes the environment of the systems called ‘terrorism’, ‘industry’, and ‘politics’? The neat idea of separating system and environment breaks down, because of the boundary problem: where does any particular social system begin and end?

This highlights a fundamental difference between simulated and real systems. In simulations, the model of the environment underlying an agent’s behaviour may be the same model used to calculate the environment’s response to the agent. Neither may be ‘correct’. Also agents may be given perfect knowledge about the simulated environment, which could be artificial.

Again, robotics seems to exist between relatively simple physical systems and more complex social systems, lying between the simulations of physical systems that can be trusted, and the simulations of social systems, many of which cannot be trusted.

Robot simulation and prediction:

As discussed earlier, there are classes of systems which are excluded from the test-by-experiment approach, for logical, moral, risk, political, finance, time, space, population extinction, contamination, and initial state replicability reasons.

With robot systems it is easier to approach replicable initial states than social systems. Although they are very sensitive to initial conditions, this can be investigated by careful laboratory experiments, run many times to obtain distributions of outcomes. In some complex systems this sets the limit of prediction.

Robotics is one of the few examples of complex system that can both be simulated and built promising to give significant insights into the ‘can you trust it?’ problem. Robot simulations are predictions of the real system. Thus robotics is a rare example of a class of complex systems whose theory and simulated predictions can be tested experimentally.

4.3 Designing, controlling, and managing robots

Robot soccer players are autonomous agents, and multi-agent teams of robot soccer players are complex adaptive systems, evolving their own self-organising behaviour, and co-evolving through playing with other robot teams. Designing such systems can give insights into general issues related to synthetic systems.

Bodies, sensors and the boundary problem

In many systems there is difficulty in knowing where the boundaries lie. For example, where does the London road traffic system begin and end? Where is the boundary of a multinational company? Where is the boundary of a human?

Robots have embodied intelligence, in which some of the behaviour of robots emerges from the physical construction of the machine, rather than from symbolic information processing [15]. The physical boundary may communicate with the environment by pushing or being pushed.

In contrast, A robot's information universe is *spanned* by its sensors, which define an *information boundary* to the system. Whereas this usually has fixed geometry for a single robot, multi-robot systems have information boundaries with more general geometrical and topological properties. Thus laboratory studies of multi-robot systems can give useful insights into the nature of boundaries of more general complex systems.

Robots as examples of synthetic systems

Synthetic systems are those that are man made, *e.g.* aeroplanes, buildings, markets, companies, medicines, computers, schools, hospitals, wars, and movies. And robots. Most synthetic systems are *designed*, by a process that involves deciding what is wanted, generating candidate solutions to the problem of producing what is wanted, and evaluating them. If these candidates fail the evaluation, new candidate solutions are generated, and the design cycle starts over again and continues until a satisfactory solution is found. The specification for designs often changes during the design process, with the addition or removal of constraints. Thus there is a coevolution between the statement of the design problem and candidate solutions until an acceptable problem-solution pair is found.

For many synthetic systems there is no *a priori* science, and the designer creates a system-specific science enabling the particular designed system to emerge from the requirements-solution coevolution.

Herbert Simon [20] writes that artificial things: 1. are synthesized by man; 2. may imitate appearances in natural things while lacking, in one or many respects, the reality of the latter; 3. can be characterised in terms of functions, goals, adaptation; and 4. are often discussed, particularly when they are being designed, in terms of imperatives as well as descriptives. This marks a distinction between designed synthetic systems and natural systems. Not only are they studied for what they are, but they are studied for what they *ought* to be. For example, a robot soccer team ought to score more goals than the opposing team.

In the laboratory of robotics, the design of robots and robot communities may inform issues to do with synthetic systems and sciences of the artificial [20].

Autopoiesis in robot systems.

The theory of autopoiesis developed by Maturana and Varela involves living systems that are endogenously controlled and self-organising. Autopoiesis means ‘self-production’, or the process by which an organism produces itself. The term autopoiesis has been applied to social systems, but this deviates from the original definition. At social level there are: self-organising systems that arise spontaneously as specific states or sequences of states, due to certain initial and limiting conditions; self-maintaining systems that self-organise to produce each other in an operationally closed way; and self-referential systems that organise the states of their components in operationally closed ways.

Autopoiesis in biological and social systems is directly relevant to robot and robot communities, as are artificial intelligence, artificial life, machine reasoning, machine problem solving - areas in which robotics has evolved. Since everything is ‘inside’ the system, robots also involve distributed data and information processing. These notions can be investigated through robotics, both at the level of the individual robot (corresponding to a biological organism), and communities of robots (corresponding to social institutions). They may become much more relevant if robot architectures move to ‘wetware’.

Language, Communication, and Culture in Robots

The building or evolution of appropriate languages is a fundamental issue in complexity science. Robot systems can be used to investigate automatic language formation, individual and group learning, as in the ‘talking heads’ experiments [16, 17, 18, 19]. Also robot systems can be used to investigate artificial cultures, as constructs emerge and become integrated into individual and group vocabularies.

Robot-developed languages lead to issues around the evolution of culture which could generalise to human systems. Robots may have the same issues of self and society as humans. Ultimately a robot can only know that which can be sensed or received as outside information, based on its own ability to process that information, in the context of information designed into it at the outset.

In principle, each robot ‘sees’ the world in its own way, leading to issues of how they can communicate through shared language. Will robots will use the same symbol to represent the same sensed entity, and does the use of the same symbolic representation imply that the same entity is sensed.

In robot soccer, a culture can emerge between robots, in which shared language can underlie very efficient communication. For example, soccer players communicate through vision, with players ‘reading’ the field, including abstracting the motives of other robots by the way they move [8].

Apart from issues of communication, there are issues of robots developing personalities. Robots may be egotistic, believing that they should keep the ball because they are most likely to score a goal. Such robots might receive the disapproval of other robots, and modify their behaviour [12]. This illustrates that robots can be used to investigate the nature of societies, and the co-evolution between agents that may cause individual robots to modify their behaviour, and possible change the behaviour of the group as a result.

5 Conclusion

The argument in this paper can be summarised as follows:

- societies of robots are complex systems: robot systems
 - exhibit emergence, self-organisation, evolution, co-evolution
 - are computationally irreducible and chaotic systems
 - can be multi-agent systems
 - require explicit symbolic and analogue representations and languages
 - require vision and pattern recognition from data grounded in sensors
 - are autopoietic systems
 - can have embodied behaviours emerging from interaction with the environment
- robot systems can be studied without the complications of human agents
 - robots don't (need not) deliberately disobey given by humans
 - robots are not (need not be) stakeholders
 - robots can be studied without the complication of observing human systems
 - robots may abstract constructs and languages not imposed directly by humans
- robot systems can be used for repeatable systematic laboratory experiments
 - practical experimental exclusions do not apply: logical, temporal, etc.
 - ethical issues do not apply to robot experiments: moral, political, etc.
 - experiments can be designed and systematically executed
- robot systems can be used to research the bottleneck problems in complexity
 - robots can be used to investigate explicit representation inside computers
 - robots can be used to investigate lattice hierarchical representation structures
 - robots can be used to investigate hierarchies of authority and control
 - robots can be used to investigate the system boundary problem
 - robots can be used to investigate automatic construct and language abstraction
 - robots can be used to investigate grounded theories
 - robots can be used to investigate emergence and automatic pattern recognition
- robots systems can be used to research fundamental issues in simulation
 - robots can be used to investigate the real-invented universe problem
 - robots can be used to investigate the 'can you trust it?' problem
 - robots can be used to investigate natural limits to prediction and what is knowable.

The questions raised in the paper are bottlenecks holding back the development of complexity science. It has been argued that they can be studied using teams of interacting robots, and that there are tremendous advantages in robot experiments, especially taking humans out of the loop and the possibility of systematic repeatable experiments. Therefore, the main conclusion to be drawn from this paper is that some of the major advances in complexity science over the next decade will be made through robotics.

References

1. Arthur, B., Increasing returns and path dependence in the economy., University of Michigan Press (Ann Arbor), 2000.
2. Bourguine, P., Johnson, J., ‘ONCE-CS: Open Network of Excellence in Complex Systems’ Coordination Action proposal under FET Open, European Community, 13-Sept-04
3. Casti, J., ‘Can you trust it?: on the reliability of computer simulation and the validity of models’, *Complexity*, 2(5), 8-11, May/June 1997.
4. Dreyfus, H. L., What Computers Still Can’t Do: A Critique of Artificial Reason, MIT Press (Cambridge, Mass), 1992
5. Gould, P., Johnson, J., Chapman, G., The structure of television, PION (London), 1984.
6. Johnson, J., ‘A language of structure in the science of complexity’, *Complexity*, 1(3),22-29, 1985.
7. Johnson, J. H., ‘Robot football, artificial life, and complexity’, *Proc. Third Int Symp. on Artificial Life Robotics*, M. Sugisaka (ed.), Oita, Japan, 19-21 Jan 1998.
8. Johnson, J. H., ‘Visual communication in swarms of intelligent robot agents’, *Artificial Life and Robotics*, 5, 1-9, 2001.
9. Johnson, J.H., ‘The “Can You Trust It?” Problem of Simulation Science in the Design of Socio-technical Systems’, *Complexity*, 6(2), 34-40, 2001
10. Johnson, J.H., Price, B.A., ‘Complexity science and representation in robot soccer’, *Proc. RoboCup*, Padova, Italy, July 2003.
11. Johnson, J.H., Rose, V., ‘Automatic construct abstraction in machine vision’, mimeo, Department of Design and Innovation, The Open University, Milton Keynes, UK. 2004.
12. Johnson and Sugisaka. ‘Egotism body language, communication, and culture in swarm robots’. ICSE2000, Coventry, UK.
13. Kitano, H., ‘RoboCup: A Challenge AI Problem’, *AI Magazine*, Spring, 1997.
14. Marr, D., *Vision*, W. H. Freeman and Co (New York), 1982.
15. Pfeifer, R. Scheier, C., *Understanding Intelligence*, MIT Press (Cambridge, Mass) , 2001
16. Solé, R., Goodwin, B., *Signs of life: how complexity pervades biology*, Basic Books, Perseus Books Group (New York), 2000.
17. Steels, L., ‘Language as a complex adaptive system’, *Lecture notes in Computer Science. Parallel Problem Solving from Nature - PPSN-VI*. Schoenaur et al, Springer Verlag (Berlin), 2000.
18. Steels, L., ‘Intelligence with representation’, *Phil. Trans. R. Soc. London. A* (2003) 361, 2381-2395, 2003
19. Steels, L., ‘Evolving grounded communication for robots’, *Trends in Cognitive Sciences*, 7(7), July 2003
20. Simon, H., *The Sciences of the artificial*, MIT Press (Cambridge, Mass), 1969.