Published as:
Dautenhan K, McOwan P W & Warwick K
'Robot neuroscience; a cybernetics approach'
in 'Neuromorphic systems. Engineering silicon from neurobiology' Editors L Smith & A Hamilton, World Scientific 113-125 (1998)

ROBOT NEUROSCIENCE - A CYBERNETICS APPROACH

KERSTIN DAUTENHAHN, PETER McOWAN & KEVIN WARWICK
Department of Cybernetics, University of Reading, Whiteknights, Reading, RG6 6AY
E-mail: K.Dautenhahn@cyber.reading.ac.uk
http://cyber.rdg.ac.uk

The sensory, actuator and neural processes of a robot are seen to be best represented and implemented as an overall system. Each is dependent on the other, and by viewing them collectively the result is greater than the sum of the parts. With carbon-based lifeforms inter-dependency of the process has resulted as an evolutionary feature. In this chapter the authors present a concept of inter-dependency, indicating a distinct approach in the construction of real world "intelligent" robots. A range of experimental material is also included to back up the discussion, based on a number of robots constructed. Details of wheeled and legged robots are given in which a key emphasis has been to integrate sensory, actuator and neural processes.

1 Introduction

Robotics provides an ideal environment in which to experiment with the three key aspects, namely sensors, actuators and intelligent decision making mechanisms. It is useful to compare robots with carbon based lifeforms, both from the point of view of related performance characteristics and also in terms of how the three areas link together and interact. This, of course, also brings into question the environment in its broadest sense. However, it is vital to reflect on how a being's (robot or carbon-based) sensors obtain a picture of the environment and what effect on the environment its actuators can have. The strengths, limits and power of the intelligence of the being are directly dependent on the sensors and the actuators, and must be taken into account in assessing the being's overall capabilities. Essentially, a being's intelligence is directly linked to the capabilities of its sensors and actuators, which in turn are linked to the environment. In this chapter we will describe how the application of the principles of cybernetics, namely the examination of full, interacting systems and the information processing strategies within, give a useful insight into robot neuroscience.

2 Walking Robots

In the field of legged robots, research has been going on for some time, the work by Raibert and Sutherland [1] being a good example. Much of this research however has been concerned with either the mechanical aspects of the legs themselves or in obtaining a reasonably accurate mathematical description of the walking process itself. The quest for autonomy in such a robot has often been overlooked in order to ensure sufficient off-robot power and processing capabilities simply for the walking action, as in Chiel at al.[2] The whole field was however thrown into disarray at the end of 1996 when Honda unveiled their P2 humanoid robot which can operate autonomously, and is capable of walking up and down stairs in a very human-like way.

As far as walking robots are concerned, the open field for research has shifted clearly towards the intelligence of the robot itself and away from the inherent walking action. Essentially the walking robot is merely a base for sensors and intelligence, the fact that the robot walks is no longer of great significance, though this movement is affected by the environment. With regard to intelligence itself, the most significant body of research is perhaps that carried out at MIT, this being nicely summarised by Ferrell[3]. One of MIT's most recent robots, Hannibal, is a completely autonomous hexapod robot, with 19 actuators, more than 60 sensors and 8 on-board computers.

Walter, Elma and Sly[4],see Fig. 1, are hexapod robots developed at Reading University. They are, like Hannibal, completely autonomous, but are less sensor ridden. Importantly their intelligent processing is carried out in a distributed fashion, by means of LON technology. This method has been directly likened to the nervous system of a cockroach in that a central (head) unit makes decisions on direction and gait, sending signals down to each of the legs. Each leg then takes care of its own actions, however should a failure occur somewhere in the robot, each leg has sufficient intelligence to continue its own operation by inferring system states. So, as with insects, if one leg fails to operate, the others can compensate automatically, in order to arrive at a successful (limping) gait.

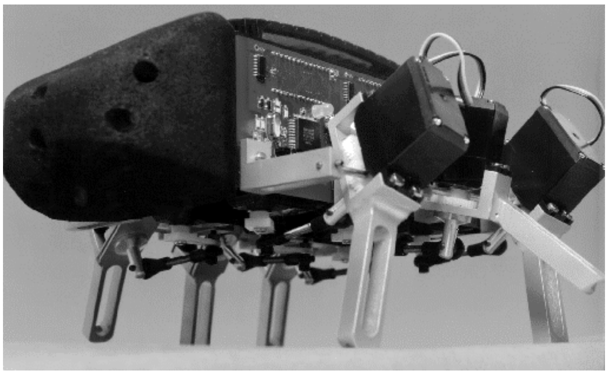

Fig.1 Elma the hexapod robot developed at the University of Reading

The three Reading robots all have ultrasonic arrays to the front left and front right of their heads, with Elma also having arrays centrally positioned. In this way they can detect the presence or absence of objects in those directions. Very little processing is then required in order for the robot to decide what to do in the event of an object being closely sited, i.e. the robot could decide either to track or move away from the object. Through infrared information the robots can have communicated information passed to them, this can be for a wide variety of uses, even switching the robot on and off!

Although the hexapod robots can only achieve a few of the lower levels of capabilities as discussed by Brooks[5], both Elma and Sly have learning capabilities. These can be used to enable the robot to learn what do with its legs in order to either track objects or to avoid bumping into objects. Such a learning experiment takes rather a long time however and at present is not always achievable given the approximate 2 hour charged battery lifetime. Practically though both Elma and Sly have the capability to learn how to position and move their legs in order to walk. This involves keeping their body off the ground, organising each individual leg and co-ordinating the legs as a whole.

Experiments are now being conducted to look into different learning strategies and the incorporation of further sensors. Interestingly both Elma and Sly have an approximate neuron equivalent of 100 cells. With this they are able to deal with fairly simple sensor information, to control a number of actuators (legs) and to learn how to co-ordinate their actuators in order to achieve the goals of walking and avoiding obstacles.

3 Wheeled Robots

In the field of mobile robotics, for a number of reasons, wheeled robots have received far more attention than legged robots. Some of the main reasons can readily be seen to be ease of actuator control, stability, co-ordination of movement,(usually much simplified), smooth power demands and higher obtainable speeds. It is also apparent that the intelligent processing requirements are not necessarily great unless fairly complex sensors are installed.

One trap not to fall into in designing such robots is to try to get the robot to look at the world in a human-like way when its sensors are not up to the job. For example "any motion requires a model of the local environment", Chatila 1995 [6], can push effort into trying to get the robot into modelling its environment. Is this necessary? How much intelligence does it need to form an environment? Do insects have much of a model of their environment? Even with animals, are all motions dependent on a model of the environment?

It is very easy in trying to go too far, too quickly, to get sensor capabilities and intelligence capabilities out of balance. Coupled with this, the extent to which the robot can affect the environment must be remembered. So a robot need not know if an object is square or round unless it is going to get close to it or interact with it in some way. The whole balance of the three elements, intelligence, sensors and actuators is clearly very important.

By taking more of a bottom-up approach, a balance can be more easily retained. The idea being to gradually build up, adding more sensors, more intelligence and more actuators whilst retaining a reasonable equilibrium. By gradually adding more and more small enhancements robots can be generally evolved in a fairly steady way.

3.1 The Reading University insects

Most work in the area of autonomous wheeled robots has concentrated on simulation studies, however a number of real experimental projects have been carried out, in particular with interacting robots, e.g. Matovic et al., 1995[7]. At Reading University reinforcement learning in wheeled robots for low level, reactive control has been an interesting target, Mitchell et al., 1994[8]. These robots, see Fig. 2, have ultrasonic sensors to the front, in the same way as the legged robots, a caster below their front and two separately driven wheels to the left and right rear.

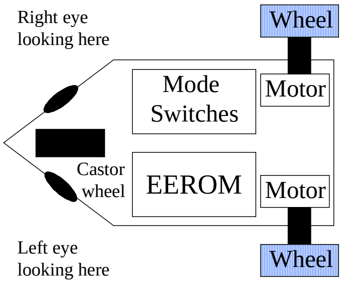

Fig.2 Schematic of the Reading university "insects"

These robots can be programmed to move around in their environment, either without hitting anything or by closely tracking a moving object. However, they can also learn a suitable strategy by means of a trial and error process. The problem can be split up into two parts, at any instant, (a) what situation is the robot in - as sensed by the robot itself, and (b) what should it do with its wheels in that situation. So, if the robot is given an overall goal of moving around but not hitting anything in its way, it can learn to satisfy the goal as follows. (a) For any given situation the robot can carry out one of a number of possible actions with its wheels, e.g. left wheel forward fast, right wheel backward slow. (b) Each possible action is described as an automata and for any situation they are associated with a fuzzy value. (c) Initially the values are randomised such that when the robot first gets into a particular situation it will try out the automata which is selected by means of a weighted roulette wheel procedure, i.e. those automata with the highest fuzzy values are more likely to occur. (d) If the automata is successful its fuzzy weighting, for that situation only, is increased with others decreased. The exact procedure is described in more detail in Mitchell et al. ,1994[8].

The situation that the robot is in at any time can also be learnt, this time by means of a Hopfield type network, see e.g. Irwin et al., 1995[9]. By means of such a network the robot can form a situational map directly dependent on its sensory input. If an object is closely positioned to the right of the robot this will most likely be witnessed through the right hand ultrasonics, possibly not through those on the left and maybe a little by those in front. This situation will be directly linked with its own set of fuzzy automata, the adapting of which was just described.

By moving around in its environment the robot can encounter different situations and by fuzzy automata based on trial and error, learn what to do in those situations. Usually, after a few minutes, the robot can work out a reasonably successful behavioural make up. Subsequently the robot should be able to move around in its environment without hitting obstacles in its own particular behavioural fashion. Exactly how it behaves is dependent not only on itself, its sensors and its actuators (wheels) but also on the environment. It is, therefore, critical when it encounters certain situations, what it does in those situations in the first instances and whether that is successful or not. The different behaviours obtainable and a fuller description of results achieved can be found in Warwick, 1997[10].

In neuron terms it has been estimated that the wheeled robots described use a mere 40 to 50 cells in order to achieve the behavioural, responsive learning described. It is worth remembering that the sensory input is relatively simple as is the required actuator signals. The system as a whole operates as a feedback mechanism, including environmental effects. For this overall process 50 cells are about right, not too many, not too few. However, even given this modest neural system, the insects begin to exhibit behaviours attributable to living systems.

4 Artificial Life experiments on interaction and dynamic coupling

Artificial Life (AL) is known as “the study of man-made systems that exhibit behaviours characteristic of natural living systems” [11]. The general methodology to synthesise life-like artefacts is to use natural and artificial systems as part of a "comparative study" [12]. The main point in the "artificial life roots of artificial intelligence"[13] is the bottom-up approach and the concept of "embodiment". The following sections discuss two robotic experiments which we performed in order to study embodied systems which are embedded in an ecological (and social) environment. We describe the scenario and control philosophy behind the experiments. Details and results can be found in Dautenhahn, 1997[14].

4.1 Example I: Dynamic Robot-Environment Couplings on a Seesaw

Keeping balance is a difficult problem for a robot, not only for a walking robot (see section 2) but also for a wheeled robot. The robot is in danger of overturning as soon as it is no longer running on flat, level ground. Moreover, its movement characteristics (speed, acceleration, in combination with a mass centre point far above the ground) can be such that it has to be able to control its body axis. In our experiments we used a hilly landscape, the "Hügellandschaft" [15] and carried out design studies on robots moving in plastic pipes (as part of a project on sewage robots in the AI-Lab at GMD, Sankt Augustin, Germany). One important common property of these two different experimental environments (ecosystems) is the problem of keeping balance. We tackle the problem of balance with a behaviour-oriented control approach.

We decided to build a seesaw as a dynamic and interactive environment. The primary motivation was not to find an efficient algorithm which is better than existing solutions in control theory, the goals for this endeavour were rather (a) Controlling a balancing robot by using a behaviour-oriented design approach. (b) Finding a cheap solution to a transportable robot-environment system which can be used for demonstration purposes. (c) Providing the system with an interactive aspect which allows humans to interact with the system and so explore the characteristics of the system.

Interactivity has two aspects in this experiment, the seesaw immediately reacts to the robot's movements, which in turn influences the movements of the robot and humans, too, can interact with the seesaw and therefore manipulate the robot-environment interactions. We therefore use the seesaw experiments[16] in order to discuss the role of the human observer and designer as an active embodied agent who is biased towards interpreting the world in terms of intentionality and explanation.

The robots are not balancing themselves in a stable environment rather they have to maintain a certain relationship to their dynamically changing environment. This approach towards constructing special-purpose-environments is supplementary to lines of research which focus on how adequate designs for robot morphologies could be chosen or evolved in a given environment[17].

The motivation to use an ecological approach, i.e. studying this issue in a co-adapted robot-environment system, originates in the assumption that robot research could learn from natural systems (animals and plants) which by their morphology and behaviour remind us of their adaptation (in ontogeny as well as phylogeny) to biotic and abiotic environmental parameters which, in an ongoing process, shape their “brains and bodies”.

The experiments were performed with small robots, which were built using *fischertechnik* components. They served as research models in order to implement and test principles of robot design and control. The robots had two driven front wheels, contact switches (binary inputs) and analogue tilt sensors. In order to have a very simple “sensitive body surface”, each robot had a belt around its body which is attached to contact sensors. The robots are at maximum 35 cm long and 35 cm wide and equipped with on-board energy supply and a special on-board computer, the “sensory-motor-brick” (developed at the VUB-AI Lab in Brussels). We used two robots with slightly different shapes each weighing 2kg.

The seesaw consists of a wooden plate and one or two supporting plastic hemispheres, so that the seesaw can change its orientation with one or two degrees of freedom. Seesaws with different tilting characteristics could therefore easily be constructed. The hemispheres were used because they allow a smooth tilting of the seesaw. The selection of the right material for the upper part of seesaw was difficult, because the weight of the robots and the seesaw had to be coadapted. In the final implementation the sensitivity of the seesaw was very high, i.e. near the zero position slight movements of the robots caused a slight tilting of the seesaw. Previous implementations, with less sensitive tilting behaviour, impaired the possibilities of controlling the robot's balancing behaviour.

4.2 Behaviour-oriented control

The robots were controlled using the behaviour-oriented approach and a C-based programming language PDL (Process Description Language[18]). The main characteristics of PDL are the concepts of “quantities” and “processes”. The processes are mappings between the incoming stream of values of the sensor quantities and the outgoing stream of values of the actuator quantities. The processes are executed in parallel, they do not inhibit or activate each other, and without a hierarchy. The influences of the processes on the actuators are summed and executed in each PDL cycle. The processes consist of **AddValue(q,x)** statements, increasing the value of quantity **q (Value(q))** by value **x**. In this way it was possible to implement an incremental change of the quantity values. Each PDL process does not specify or trigger a specific behaviour, they rather specify how the system should change in specific situations. A set of PDL processes can, in a specific context, yield a specific behaviour. A single PDL processes can belong to different sets of “behaviours”.

Behaviours (generally considered as observable agent-environment regularities) are not part of the control program, they are only observable when the system is put into a specific real-world context. The PDL approach is an alternative to the subsumption architecture which defines behaviours and relations (e.g. inhibition) among them[5].

The balancing problem is approached by using a hill-climbing strategy by defining two processes: turning-on-the-spot (turnRobot) and translation (climbRobot). Note that there is no central process called “balance”. Trajectory traces documented the progress of the experiments. The orientation of the robot's body axis is measured with analogue inclination sensors, fixed on the robot's chassis in rectangular orientation to each other. Experiments were conducted with one and two degrees of freedom and different starting positions of the robot. Additionally, a second robot was used which could be guided by a lightsource (phototaxis). The experiment then showed how the balancing robot reacted to disturbances caused by the second robot which was running on the seesaw. The light source could be moved by a human who could therefore influence the behaviour and dynamics of the robot-seesaw system.

Experiments with one and two degrees of freedom revealed some interesting dynamic effects, such as oscillation, when (similar to a resonance effect), the seesaw and robot changed direction with similar frequency. For a few seconds, when both movements were in phase, the amplitude of the seesaw increased until slight shifts occurred and the robot-environment system went out of phase. The robot had no means of detecting such an (from an observer point of view) interesting situation in terms of dynamics. For future investigations it might be interesting to have a robot which is able to control and evoke deliberately such situations, like a child sitting on a swing and trying to increase the amplitude.

This section describes how the balancing problem can be solved by exploiting the dynamic robot-environment interactions and using a parallel, behaviour-oriented control architecture. The global pattern of “balancing behaviour” (from an observer point of view) resulted from the interactions of a few processes which only used information about the current position of the robot's body axis.

4.3 Example II: design of a helping scenario between heterogeneous robots

The task was to implement robot “social” interactions in a “helping scenario”, given two autonomous, heterogeneous robots. The helping behaviour experiment can be described as follows: The idea was to implement one robot (equipped with light sensors) which works as the “seeing-eye robot” for a second robot which does not have light sensors. An ecological context is used, i.e. a light source is mounted on top of a charging station where the robots can recharge their batteries. Thus, guiding a robot to a light source means helping to find "food" (energy) in order to survive. Additionally, the seeing-eye robot has to recognise the "blind robot" in order to be motivated to go to the charging station. Thus the seeing-eye robot does not automatically guide any other robot to the charging station, communication and recognition of an individual has to take place. The guide recognises the follower by a specific spatio-temporal pattern of physical contact: the follower has to perform a full circle around the guide, sequentially touching the bumpers located in the ring of contact sensors. Thus, we implemented recognition on a purely behavioural level without the need of a recognition module in the control architecture.

The robots ran in the hilly landscape, the basic technology of the robots was similar to the one presented in the previous section. Fig. 3 show the seeing-eye robot (left) and the follower in a typical keeping contact situation. We used PDL to implement the helping scenario in both guide and follower robot. A few "basic behaviours" were defined as follows:

(1) *Obstacle avoidance.* This behaviour is implemented by 16 processes, each one reading the sensor signal from one of the 16 bumpers. Obstacle avoidance is done by an alignment procedure. Each process consists of two addvalue statements: **AddValue(MotorRight, rt-right x Value (ButtonN))** and **AddValue(MotorLeft, rt-left x Value (ButtonN)). rt-left** and **rt-right** are constant in each process, they specify the rotation tendency given to the left or right motor. The general idea for this kind of obstacle avoidance behaviour is that the robot turns until the obstacle is located at its side.

(2) *Phototaxis.* This process consists of two addvalue statements: **AddValue(MotorRight, (Value(EyeLeft) - Value(EyeRight)) / 5)** and **AddValue(MotorLeft, (Value(EyeLeft) -Value(EyeRight)) / 5))**

(3) *HillAvoidance.* The inclination sensor measuring the forward-backward axis (quantity Horr) was used as follows: **AddValue (MotorRight, Value(Horr) - 102)** and **AddValue(MotorLeft, 102 - Value(Horr))**

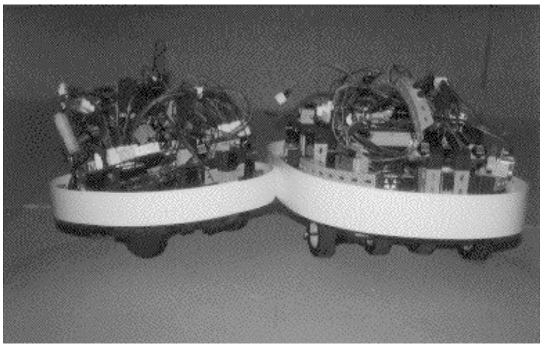

Fig. 3 The "helping behaviour" experiment

We did not use any conditions for the processes. Thus, for the implementation of the “helping scenario” the processes had to be carefully designed, e.g. for triggering specific processes due to specific sensor readings we had to reduce all and/or conditions to multiplication/division and addition/subtraction. A number of internal quantities had been used to control the set of processes, comparable to a “motivational system”. The system was fully hand-crafted, e.g. the constants in the PDL processes (see above) had been set by systematic tests. The experiments were documented with a video camera. Two of the main problems which occurred were: (a) Since none of the robots used distance sensors the communication and following process was very brittle, i.e. susceptible to disturbances, so that the robots lost contact (due to the hilly structure of the ground, friction, inertia etc.). (b) The robots acted purely reactively and locally. They did not differentiate between contact to a wall or to another robot. Thus, when the follower touched a wall while following the guide it got distracted by the additional sensory inputs.

The experiment worked best on a plain ground without any disturbances. In that case the scenario worked quite well so that from the point of view of people who observed the scenario a “social story” was “enacted”. In these cases the big gap became obvious between the human observer, who is biased to interpret the world in terms of intentionality and who easily attributes elements of his/her theory of social cognition (we discuss this issue in more detail in Dautenhahn, 1997[16]), and the technological basis employed, in this case a simple reactive, parallel control architecture.

4.4 Summary of Artificial Life robot interaction experiments

The robots were situated, since they completely depended on on-line, real world sensor data which were used directly in a behaviour-oriented control architecture. We did not use any simulation. The robots did not utilise any world model. The robots were embedded, since robot and environment (social and non-social, e.g. the seesaw and another robot) were considered as one system, e.g. design and dynamic behaviour had to be carefully co-adapted. However, in comparison to natural living systems the robots have a “weak” status of embodiment. For example the body of the robot is static, the position and characteristics of the sensors and actuators are modified and adapted to the environment by hand, not by genuine development.

5 Conclusions

In this chapter the important issue of a balanced approach to robot neuroscience has been described. It has been stressed that the neural aspects of a robot should not be designed in stand-alone fashion but rather the design process should also take into account the sensory and actuator requirements. A study of the required complexity of the robots sensory and motor interaction with the environment is an essential part of the developing an understanding of these artificial systems. Examples of robot constructions of both legged and wheeled varieties were given to supplement the thesis presented. In section 4 we presented two examples of physical systems which were dynamically coupled to its social and non-social environment by using a behaviour-oriented control approach. The experiments point the way toward building more advanced and “truly” embodied systems, in the sense which is for instance described in Dautenhahn 1997[16]. Thus, by perusing a course of robotics research drawing on all the core themes of cybernetics, we may develop and understand ever more “life-like” creations.

Acknowledgements

The authors would like to acknowledge the following people. Dr David Keating and the mobile robotic research group at the University of Reading. The robotics experiments described in section 4 were supported by the German National Research Center for Information Technology (GMD), Sankt Augustin, Germany and the University of Bonn, Germany.

References

1. Raibert M. & Sutherland I.E. (1983) Machines that walk, *Scientific American*, Vol. **248**, pp. 32-41.
2. Chiel H, J., Beer R.D., Quinn R.D. & Espechied K.S. (1992) Robustness of a distributed neural network controller for locomotion in a hexapod robot, *IEEE Trans. on Robotics and Automation*, Vol. **8**, pp. 293-303.
3. Ferrell C. (1995) A comparison of three insect inspired locomotion controllers, *Robotics & Autonomous Systems*, Vol. **16**, pp. 135-159.
4. Wenn D., Glover N. & Warwick K. (1995) WALTER an autonomous walking LON controlled robot, *LON users Conference* San Francisco 23-25 May 1995.
5. Brooks R.A. (1986) A robust layered control system for a mobile robot., *IEEE Journal of Robotics and Automation* **2**, pp.14-23.
6. Chatila R. (1995) Deliberation and reactivity in autonomous mobile robots, *Robotics and Autonomous Systems*, Vol. **16**, No. 2-4, pp. 197-211.
7. Matavic M.J., Nilsson M. & Simsarian K.T. (1995) Co-operative multi robot box pushing, *Proc. IROS-95*, Pittsburgh.
8. Mitchell R.J., Keating D.A & Kambhampati C. (1994) Learning strategy for a simple robot insect, *Proc. Int. Conference Control 94*, Warwick, pp. 492-497.

9. Irwin G., Warwick K. & Hunt K. J.(eds) (1995) Neural network applications in control, *IEE Press*.

10. Warwick K. (1995) March of the Machines, *Century*.

11. Langton C. G. (1989) Artificial Life. *Proc. of an Interdisciplinary Workshop on the Synthesis and Simulation of Living Systems*, Los Alamos, New Mexico, September 1987,pp. 1-47.

12. Taylor C. & Jefferson D. (1994) Artificial life as a tool for biological inquiry. *Artificial Life* **1**, pp. 1-13.

13. Steels L. (1994) The artificial life roots of artificial intelligence. *Artificial Life Journal* **1**

14. Dautenhahn K. (1997) Biologically inspired robotic experiments on interaction and dynamic agent-environment couplings, *Proc. Workshop SOAVE'97*, Selbstorganization von Adaptivem Verhalten, Ilmenau, 23.-24. September 1997.

15. Dautenhahn K. (1995) Getting to know each other - artificial social intelligence for autonomous robots. *Robotics and Autonomous Systems* **16**, pp. 333-356.

16. Dautenhahn K. (1997) I could be you - the phenomenological dimension of social understanding. *Cybernetics and Systems* **25**, pp. 417-453.

17. Sims K. (1995) Evolving 3D morphology and behavior by competition. *Artificial Life Journal* **1**, pp. 353-372.

18. Steels L. (1992) The PDL reference manual. Memo published at the Artificial Intelligence Laboratory, Vrije Universiteit Brussel, March 1992.