

Active Perception, Navigation, Homing, and Grasping: An Autonomous Perspective

Dario Floreano
Laboratory of Cognitive Technology
AREA Science Park
Trieste, Italy
dario@psicosun.univ.trieste.it

Francesco Mondada
Laboratory of Microcomputing (LAMI)
Swiss Federal Institute of Technology
Lausanne, Switzerland
mondada@di.epfl.ch

Abstract

Perception is needed for action, not for the pure sake of the construction of abstract representations, although it does not exclude the role of internal representations for mediating complex behaviours. We think that, for the purpose of building autonomous robots, active perception requires specific recipes for three related aspects: the design of the physical sensory system, the modality and type of information extracted, and the structure and functioning of the control system. We outline a set of solutions for these three aspects and describe their implementation on a real mobile robot through a set of three different experiments using a combination of neural networks and genetic algorithms. The results show that active perception is a useful feature that is exploited by autonomous agents. The experiments show that the combination of genetic algorithms and neural networks is a feasible and fruitful technique for the development of active perception in autonomous agents.

Keywords

GENETIC ALGORITHMS – NEURAL NETWORKS – AUTONOMOUS AGENTS

1 Active Perception and Autonomous Agents

The active sampling of the visual scene plays a fundamental role in the development of sensory-motor co-ordination of human infants [GIBSON91]. The construction of a stable and reliable representation of the external world is gradually built through a long process of interaction with the environment: babies actively explore the world and their own body for several months after birth. It is reasonable to argue that exploration is equally important for most animals because, since there are not specific vantage points of view, the sensory appearance of the world does not convey a constant representation. Even in very simple ecosystem, i.e. where the animal has a limited number of interactions with a fairly simplified environment, the autonomous choice of certain aspects of the information available may turn relevant for smarter behaviours. These considerations, that have been well understood by visual scientists, become important for those researchers that are interested in building autonomous robots by pursuing a behaviour-based approach to the development of the control systems.

Generally speaking, the adjective *active* stresses the importance of the autonomous extraction of the relevant information by the agent itself, minimizing the role of an external controller. It also suggests that perception is needed for action, not for the pure sake of the construction of abstract representations, although it does not exclude the role of internal representations for mediating complex behaviours. We think that, for the purpose of building autonomous robots, active perception requires specific recipes for three related aspects: the design of the physical

sensory system, the modality and type of information extracted, and the structure and functioning of the control system. We will address these three topics in the following sub-sections. In the second part of the paper we will describe a series of experiments with a real robot where we have successfully implemented these ideas for the development of increasingly complex autonomous agents.

1.1 Active perception: physical requirements

The sensory system of animals is not a passive device for gathering information. If we look at the natural diversity of biological organisms, it becomes evident that its structure and its functioning are shaped by the individual needs of the organism in order to survive in its own environment. Echo-location, wave length, light intensity, sound-waves, tactile stimulation are some among the various sources of information exploited by biological organisms. The physical features and the layout of the sensors employed by the animals are tailored to the physical properties of their own environmental niche and to the range of possible interactions between the animal and the external world. It does not make much sense to ask whether the behavioural and cognitive abilities of the animal depend upon the complexity of its own sensory system, or whether the characteristics of the sensory system depend upon the behavioural and cognitive abilities of the animal. The question is ill-posed because there is not such a dicotomy. The sensory-motor apparatus of biological organisms has co-evolved with their nervous system to allow the animal to survive in the environment. With regard to this point, an animal is a dynamic system composed of several parts in equilibrium.

The extent in which we are interested in reproducing biological solution, both for constructing autonomous robots and for better understanding the underlying principles, these considerations have important bearings for the design of artificial autonomous agents. Traditional wisdom tells that robots should be equipped with a reliable, precise, and complex sensory apparatus. In the old and unconscious desire to achieve human performances, many engineers try to employ sophisticated video cameras on the top of the robot producing several Kbytes of information per second in the vain attempt to mimick the structure of the human visual system. Also, it has been recently stated [TOURETZKY94] that theories of rat behaviour cannot be tested on mobile robots because of the inadequacies of the sensory devices, directly implying that faster hardware and better algorithms should be coupled with colour video cameras for a more reliable low-level visual processing. We believe this approach to be misleading for two reasons. The first is that in most cases we do not need such a complex perceptual system for achieving a number of desired tasks, such as avoiding obstacles, wall-following, reaching for locations, grasping objects (although one may argue that this holds only in simplified situations). The second – more serious – reason is that if we have several Kbytes of information to process, we must choose what are the relevant parts and how to pre-process them (exactly what the retina and the first relay nodes in the mammal visual system do).

A solution to this apparent empasse is to employ simpler sensory devices that are commensurate to the range of behaviours expected from the agent and tailored to the type of interactions between the agent and its environment. If we want to perform simple, but smart navigation and obstacle avoidance, we will not need more than a bunch of non-linear and noisy infrared sensors (as we will show in our first experiment). By only adding some weak light sensors we can have a much richer behaviour with a number of emergent rat-like solutions (as described in our second experiment). The same sensors, appropriately exploited, are enough even in the case when the robot is required to grasp objects (as shown in our third experiment).

1.2 Active perception: sensory information

Active perception means an autonomous choice of the sensory information extracted from that available and of the type of pre-processing performed on it. It also implies that, in some circumstances, the agent must explore the environment in order to find the most appropriate information for performing the correct actions. Exploration is needed in two cases: when the system is developing and when the local information is not sufficient for performing an appropriate behaviour. All these features – namely, extraction of the relevant information, choice of the filter, exploration – cannot be artificially imposed from the outside, but should become emergent properties of a truly autonomous agent interacting with its environment.

The agent must also have full access to the rough information coming from its own sensors. This is not what currently happens with several models that claim to operate in biologically plausible conditions. In most cases the sensory information is pre-processed by the experimenter. In many computational models the agent is given pre-computed distances, angles, and/or compass directions from some object or landmark. Although it may be claimed that this is the sort of operation that many animals indeed perform on the rough stimulation, this choice implies a lot of non-trivial abilities and assumptions: the animal is assumed to distinguish figures from the background, to locate them in space, to exploit some gradient information for computing distances, and to employ a specific metric system. It also implies that such information is computable. But:

- This is not necessarily the type of information needed by the agent to perform the most appropriate behaviour.
- The low-level information is lost: the internal structures of the control system cannot access it anymore.
- The sort of information available may affect the type of behaviour displayed by the agent, especially if all the redundancy and richness of the primitive stimulation is cleared out. You get what you gave.
- The sensory information gathered by real agents in real worlds is essentially local. Hence, the information required to compute angles and distances may be not always available (as for computer simulations). This is one of the reasons why you would expect emergent exploration in a system that deals with rough sensory stimulation.

For all the reasons listed above, we believe that the *active perception approach* can be sustained only if the agent has full access to the primitive activation of its own sensors. During its autonomous development it will select those aspects that are most relevant for the behaviour, for example by discarding the information coming from some sensors, enhancing that coming from others, exploiting non-linearities, combining activations, and appropriately filtering it.

1.3 Active perception: control structures and learning

Such autonomy in the choice of the low-level sensory information must be supported by an adequate control system that can deal with all these micro-features. Also, since the agent is supposed to adapt its own sensory system through a dynamic interaction with the environment, the control system employed should be plastic and evolvable in some automatic fashion. We think that artificial neural networks are the best candidates at the moment. An artificial neural network can deal separately with the activation coming from every single sensor. This

information can be enhanced, transformed in inhibition or excitation, filtered, and variously combined with the information coming from other sensors. Neural networks are well-suited structures for real world operation because most architectures and neuron models display a good tolerance to noise and possess intrinsic abilities to generalize to novel stimuli: this means that the agent will be very likely to produce adequate behaviours even though the sensory information is not exactly the same as that encountered in previous circumstances, or it is partially missing. The notion of time can be easily embedded in an artificial neural network, either by using recurrent dynamic nodes [PEARLMUTTER90] or by resorting to recurrent connection between nodes.

There are several techniques to change the synaptic weights of a neural network. In case of autonomous agents it is rather difficult to employ strictly supervised learning algorithms (such as Backpropagation, Boltzmann Machine, or Learning Vector Quantization) because the correct output is not always available or computable. We have employed genetic algorithms [GOLDBERG89] because they do not require a direct specification of the desired values for the units in the network, that is they do not necessarily specify what specific behaviours must be used by the agent to solve a given problem. This turns to be an interesting features in complex environments where the agent may take several different and equally valid decisions at any given moment, each of them leading to a different sequence of actions and experiences which affect the final behaviour outcome. This aspect is important for the autonomous development of control systems. The extent to which the experimenter controls the behaviour of the individual with genetic algorithms is given by the fitness function. As we will show in our second experiment, the emergence of complex and ecologically-grounded behaviours does not necessarily require detailed fitness functions.

2 The Experiments

We have performed three different experiments by resorting to the principles of free interaction and active perception outlined above. In all the experiments the evolutionary procedure has been carried out entirely on a real miniature mobile robot – Khepera – designed and built at the LAMI.

2.1 Experimental setup

Khepera is a miniature mobile robot of circular shape featuring 55 mm of diameter, 30 mm of height, and 70 g of weight [MONDADA93] (Figure 1). It is supported by two wheels, each controlled by a DC motor with an incremental encoder (12 pulses per mm of robot displacement). Khepera design and mechanical solutions are well-suited for experiments with autonomous agents. In its basic configuration it has 8 infra-red proximity sensors placed around its body (6 on one side, 2 on the other) that can also be enabled to measure ambient light. Khepera has a modular structure: various modules and sensors can be added without need to change the software. This allows the experimenter to employ the most suitable configuration for the type of environment and behaviour under study. It also affords the addition of further devices during the development of the control structure, thus mimicking long term adaptation in Nature where there is a co-development of the nervous system and of the physical characteristics of the organism. Khepera can be attached to a workstation by means of a light-weight aerial cable and specially designed rotating contacts. This configuration allows the experimenter to exploit the computational power (e.g., for neural network activation and on-line analyses) and




Figure 1. Khepera, the miniature mobile robot.

memory storage capabilities (for recording all the relevant data of the experiment) of the workstation, while all low-level processes (such as sensor reading, motor control, and other internal operations) run on the micro-controller that is placed on the robot. The cable can also provide electric power, an interesting feature to have when genetic algorithms (that require long periods of time) are employed.



Figure 2. A close view of the robot with the "helmet" for capturing laser signals and the laser device on the background.

In research on autonomous agents the analysis of the resulting solutions and behaviours is very important. It is often the case that the emergent solutions cannot be easily understood with traditional analytic tools [MCFARLAND92]. For this reasons, it can be useful to resort to the methodologies employed by ethologists and neurophysiologists. These techniques require a set of appropriate tools (hardware and software) for recording, storing, and analysing behavioural and internal variables of the agent while it is freely acting in its environment. For this purpose, an external positioning laser device has been built at the LAMI (Figure 2). During the adaptation process, or at a later stage, the robot can be equipped with a special "helmet" for capturing the laser light and computing exact positions. This information is not passed to the agent's neural network, but is sent through the cable to the computer where a software displays and combines behavioural data with internal network variables. This allows us to perform standard neuro-ethological analyses on the agent in its natural settings.

In all our experiments we employed a standard genetic algorithm [GOLDBERG89] with fitness scaling and biased mutations [MONTANA89]. The population size was fixed and set to 100 individuals. The neural network synaptic weights and thresholds were coded on the chromosome as floating point numbers. In all our experiments the network architecture could not be modified by the genetic algorithms. The neural networks employed were multilayer perceptrons of sigmoid units with a set of recurrent connections at the hidden layer [ELMAN90]. For the first experiment on navigation we did not employ hidden units; hence, the recurrent connections were at the output layer. Each chromosome of a population was in turn decoded into a neural network, the input units were clamped to the sensors activations and the output unit values were directly used to control the speed of each individual wheel (we did not employ action commands such as "turn x degrees to the left"); the robot was let free to move for a certain amount of time while its fitness was recorded and stored away for selective reproduction. After all the individuals in the population had been evaluated, we applied the three genetic operators (roulette-wheel selective reproduction, chromosome crossover, and gene mutation). Between one individual and the next, the robot was made to perform a random motion for 5 seconds, to avoid artificial position inheritances between neighbouring chromosomes in the population file.

2.2 Navigation and obstacle avoidance

In the first experiment we attempted to develop a very simple, well defined, and widely studied behaviour in order to test our methodology and some of our basic assumptions. The robot was put in a 50x80 maze with irregularly shaped walls and the neural network was connected to the 8 infrared proximity sensors and to the two motors. The fitness function Φ was designed to select the individuals who could keep low sensor activity and high motor activation as follows,

$$\Phi = V \left(1 - \sqrt{\Delta v} \right) (1 - i) \quad (1)$$

where V is a measure of the average rotation speed of the two wheels, Δv is the algebraic difference between the signed speed values of the wheels (positive is one direction, negative the other) transformed into positive values, and i is the activation value of the proximity sensor with the highest activity (all the three variables assume values between 0 and 1). We let the system to evolve for 100 generations (although the best individuals around the 50th generation already display optimal behaviours), until the average fitness reached a stable level. The resulting behaviour is a smooth navigation and efficient obstacle avoidance (see [FLOREANO94A] for more details and other considerations).

From the analysis of the resulting behaviours, we noticed that all the individuals in the population had developed a preferential direction of motion, corresponding to the side where the higher number of sensors are located (this happened for three out of three genetic runs, starting from initial different synaptic weights). This autonomous choice was the result of a self-selection of the most relevant sensory devices for performing adequate actions. Individuals that moved in the other direction (where there are only two sensors) disappeared soon from the population because they probably failed to accurately detect (and thus avoid) spiky corners or plain walls from certain angles.

When we compared the resulting-network behaviour with the behaviour of a Braitenberg's vehicle [BRAITENBERG84] designed for obstacle avoidance and implemented on Khepera (Figure 3), our agent performed significantly better by finding its way out of symmetric situations (such as corners or dead-ends) even though it had not been specifically designed for those situations.

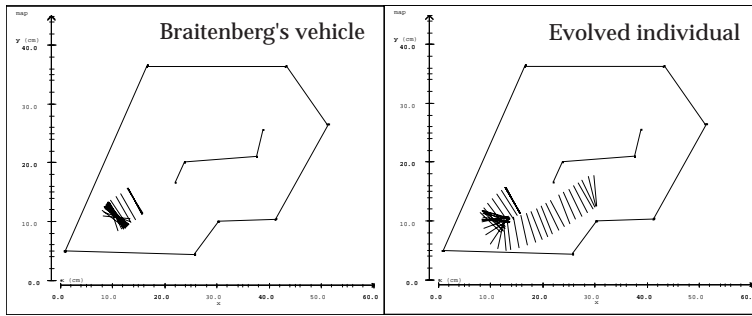


Figure 3. *Left:* Braitenberg’s vehicle, implemented as a neural network, gets trapped in a corner. Each sensor sends excitatory signals to the ipsilateral motor and inhibitory signals to the contralateral motor. Given its symmetric and feed-forward structure, this control system gets stuck when it receives equal stimulation from two opposite sensors. The small bars in the figure represent the axis between the two wheels. The position was computed and plotted every 100 ms by using the information provided by the laser device. *Right:* An evolved agent can easily negotiate symmetric sensory stimulation by exploiting the recurrent connections at the output layer. This agent never gets stuck.

2.3 Homing

The results described in the previous experiment were encouraging because the paradigm employed provided a number of interesting and efficient solutions without much human intervention. Nevertheless, several aspects of the resulting behaviour were explicitly described – even though the level of description was very different from that employed by the classic approach – in the fitness function. This was in part due to the fact that, whereas biological organisms have internal drives, our simple robot did not have particular needs to satisfy: the fitness function was an artificial mean of providing the robot a goal. In the second experiment we decided to use a more ecological setup by providing the robot with a simulated battery¹ and the environment with a restricted area where the robot could charge its battery. The environment was similar to that employed in the previous experiment (although the geometrical shape was more regular and without obstacles in it). The floor of the recharging area was painted black and a tower with several small DC lamps was placed above it illuminating the whole arena (Figure 4). The neural network received input from the eight proximity sensors, from three ambient light sensors (one placed beneath the robot surface to detect the floor brightness, and the other two on the two opposite sides of the robot), and from a sensor of battery charge. The battery had a simulated linear discharge that allowed an autonomy of approximately 20 seconds (about 50 actions). If the robot happened to pass by the recharging area the battery was automatically and instantaneously recharged. In order to avoid unlimited life durations, we established a top limit of 60 seconds per individual, after which the life was automatically truncated. The fitness function employed was a simplified version of that used in the previous experiment, without the term $(1 - \sqrt{\Delta v})$: it did not include any term for driving the robot toward the recharging area or for processing the information coming from the ambient light sensors.

Both the average fitness of the population and the fitness of the best individuals steadily

¹We could have used as well the battery already placed on Khepera, but the genetic run would have lasted more than 60 times longer.




Figure 4. The environment with the light tower, the recharging area (the floor is painted black), and the robot.

increased during 240 generations, showing that the robot learnt to move around the environment without hitting the walls (see [FLOREANO94B] for further details, full analysis, and other considerations). In the meanwhile, also the duration of the lifespan of the individuals increased and reached the top limit after 150 generations (Figure 5). This result means that the robot autonomously learnt to find the recharging area by exploiting the light gradient information. By periodically recharging the battery, the robot could perform a higher number of actions and thus obtain higher fitness scores. In order to analyze the agent's behaviour we equipped Khepera with the helmet for capturing the laser light and computing its position. We downloaded the neural network corresponding the best individual of the last generation into the robot, and synchronized the helmet device with the network cycles of activation. The information about the robot exact position, the sensor and motor state, and the activation of the internal nodes while the robot was freely moving was passed to the workstation where a special software performed neuroethological analyses. These analyses showed that the robot spent most of its time

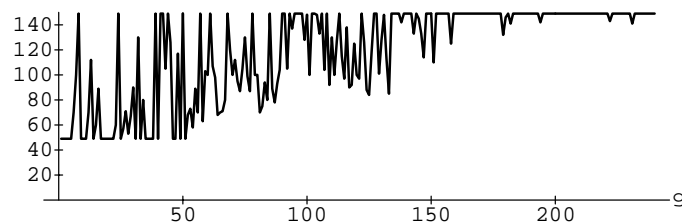


Figure 5. The number of actions performed by the best individuals at each generation. Each action lasts about 380 ms. Since each individual starts its own life with a fully charged battery, 50 actions are a lower bound (see text).

wandering around the environment without hitting the walls, but when the battery reached a minimum level, it started orienting toward the recharging area. The critical battery level

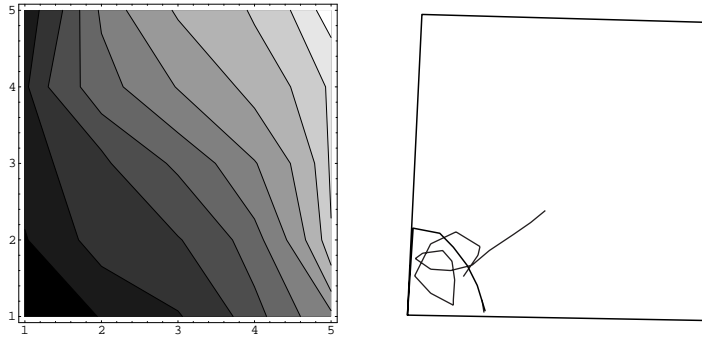


Figure 6. *Left:* Activation levels for an internal node measured at several locations evenly distributed in the environment. The recharging area is located in the bottom left corner (corresponding to the highest activation level). *Right:* The trajectory performed by the robot when the blackpainting is removed from the recharging area and the battery is not recharged. All data points have taken using the laser device. The recharging area boundaries are drawn here only for sake of clarity.

that corresponded to a change in behaviour was autonomously established by the robot and varied depending on the robot location. The robot was extremely precise in timing its return to the recharging area, spending about 90% of its energy moving in the outside environment. The accurate evaluation of the battery residual time and the location of the recharging area was allowed by an internal representation of the environment. Most of the nodes displayed activation levels that mapped the environment geometry (Figure 6), providing the information necessary for performing the most adequate actions. This internal map of the environment was autonomously developed by extracting the light gradient from the information provided by the light sensors.

In order to observe the role played by the sensors beneath the robot surface, we cleared the black paint from the recharging area and we did not recharge the battery when the robot arrived there. The robot was disturbed by this change of the environment. When it reached the area, it started to explore its surroundings without leaving that zone until its battery was completely exhausted (Figure 6). This behaviour, that resembled that of animals put in similar situations, evidenced the role played by active perception as an active search for the relevant information for survival.

2.4 Grasping

In a further attempt to make the setting even more ecological, we decided to make the fitness function simpler by eliminating the two components employed in the previous experiment. Those components partially specified what the robot should have done to achieve higher fitness, i.e. avoid obstacles and keep motors active. We thought that this behaviour specification could be lifted if a number of objects were scattered in the environment and the robot provided with a gripper module. In that case, the fitness would be only the number of objects gripped.

Trajectories, homing, wall avoidance, and other behaviours (such as searching for and, possibly, grouping the objects) would be purely emergent.

However, grasping an object is not a simple task. It involves object recognition and sensory-motor co-ordination. One of the aspects of grasping consists in a sequence of preparatory decisions and actions. These processes include the choice of the object to be grasped, deciding about the feasibility of the final act, and positioning the body system in an adequate position. In the preliminary explorations reported here, we addressed only the reaching-for-grasping task



Figure 7. The environment used for evolving the gripping abilities.

to test the feasibility of our approach. As for the experiments described above, we let the robot free to evolve the adequate sensory processing and control system during its own interaction with an arbitrary environment by employing the same genetic algorithm and the same neural network structure. The robot employed in our explorations was the miniature mobile robot Khepera with an added gripper module. The environment was an arena of the same size of that employed in the homing experiment. The objects were a number of white balls (38 mm of diameter) made of pressed paper (Figure 7). Only the activations of the proximity sensors were given to the neural network. The neural network had three output units, two controlling the motor activation, and the third the initiation of a gripping action. At this stage the gripping action itself was implemented as an automatic procedure: when the neural network decided to perform a grip, the gripper was lowered, the two segments were closed, the gripper was lifted well above the robot top surface, the segments were opened again, and finally the gripper was positioned in the resting position (Figure 7). The fitness function was simply the number of balls gripped. In a number of experiments with different initial random synaptic weight values we observed an improvement across generations of the ability to grip balls (Figure 8). We

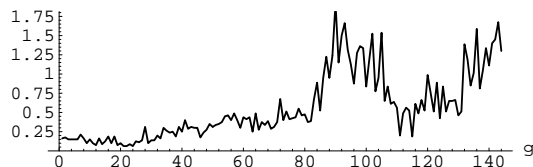


Figure 8. Average number of successful grips across generations. The temporary drop in performance around the 104th generation corresponds to the subtraction of 50% of the balls.

started by throwing 10 balls in the arena and let the robot adapt to this concentration. When its performance reached a peak level (104th generation), we subtracted 5 balls from the arena and continued evolution. After a temporary drop in performance, the agent could rapidly reach again the same performance level by improving its search strategies (see [FLOREANO92] for a similar procedure). Although it is difficult to plot typical trajectories because they depend very much upon the configuration of the balls, the strategy employed by the best individual was rather simple. The robot moved "backward" (i.e., in the direction where only two proximity sensors are placed) until it detected something; then it turned on itself until the frontal sensors corresponding to the gripper aperture detected the presence of the ball, and it activated the gripping sequence. If some of its sensors were still active, it performed a further turn on itself to position the gripper in the correct direction and performed another grip. This strategy was indeed very smart because the proximity sensors could not detect objects that were more distant than the gripper segments. It was thus impossible to approach them correctly (i.e. in the direction of the gripper aperture) from distance. The turning strategy overcame this limitation by appropriately exploiting the sensory information available.

3 Conclusion

We have showed that an approach that combines genetic algorithms and neural networks can support active perception in an autonomous robot. The results of our experiments showed that active perception is a useful feature that is exploited by autonomous agents. Choice of frontal direction, construction of internal maps by means of light gradient information, active exploration in search of missing features, and body positioning for accurate grasping were all emergent features of the active perception approach. The message conveyed by these results is that the perceptual system, both at the level of physical devices and of the control system, can autonomously sample and measure the information necessary for the appropriate action if it is let free to adapt to the dynamics of the agent-world interactions. The shift from careful design to autonomous adaptation is also a promising approach to the study of attention in simple agents. Orientation toward the recharging area only when the battery reaches a minimum level (in the second experiment), indicates that the organism selects the relevant information for adequate actions judging upon its own internal conditions.

In our current research we are trying to integrate the three experiments in a single setup (the environment features, the neural network structure, and the sensory-motor layout) where the fitness function is simply represented by the number of balls grasped. We are also trying to implement a simple form of associative learning during life (with pure Darwinian inheritance) to speed up the adaptation process in such complex situations.

Acknowledgments

We would like to thank Edo Franzi and André Guignard for the important work in the design of Khepera, and Yves Cheneval for the analysis software. This work has been supported by the Swiss National Research Foundation (project PNR23).

References

- [BRAITENBERG84] V. Braitenberg. *Vehicles. Experiments in Synthetic Psychology*. MIT Press, Cambridge, MA, 1984.
- [ELMAN90] J. L. Elman. Finding structure in time. *Cognitive Science*, 14:179–211, 1990.
- [FLOREANO92] D. Floreano. Emergence of Home-Based Foraging Strategies in Ecosystems of Neural Networks. In J. Meyer, H. L. Roitblat, and S. W. Wilson, editors, *From Animals to Animats II: Proceedings of the Second International Conference on Simulation of Adaptive Behavior*. MIT Press-Bradford Books, Cambridge, MA, 1992.
- [FLOREANO94A] D. Floreano and F. Mondada. Automatic creation of an autonomous agent: Genetic evolution of a neural-network driven robot. In D. Cliff, P. Husbands, J. Meyer, and S. W. Wilson, editors, *From Animals to Animats III: Proceedings of the Third International Conference on Simulation of Adaptive Behavior*. MIT Press-Bradford Books, Cambridge, MA, 1994.
- [FLOREANO94B] D. Floreano and F. Mondada. Autonomous and self-sufficient: emergent homing behaviours in a neural-network driven mobile robot. Technical Report R94.14I, Laboratory of Microcomputing, Swiss Federal Institute of Technology at Lausanne, May 1994.
- [GIBSON91] E. J. Gibson. *An Odissey in Learning and Perception*. MIT Press, Cambridge, MA, 1991.
- [GOLDBERG89] D. E. Goldberg. *Genetic algorithms in search, optimization and machine learning*. Addison-Wesley, Reading, MA, 1989.
- [MCFARLAND92] D. J. McFarland. Autonomy and self-sufficiency in robots. AI-Memo 92-03, Artificial Intelligence Laboratory, Vrije Universiteit Brussel, Belgium, 1992.
- [MONDADA93] F. Mondada, E. Franzi, and P. Ienne. Mobile robot miniaturization: A tool for investigation in control algorithms. In *Proceedings of the Third International Symposium on Experimental Robotics*, Kyoto, Japan, 1993.
- [MONTANA89] D. Montana and L. Davis. Training feed forward neural networks using genetic algorithms. In *Proceedings of the Eleventh International Joint Conference on Artificial Intelligence*, San Mateo, CA, 1989. Morgan Kaufmann.
- [PEARLMUTTER90] B. A. Pearlmutter. Dynamic Recurrent Neural Networks. Technical Report CMU-CS-90-196, School of Computer Science, Carnegie Mellon University, December 1990.
- [TOURETZKY94] D. S. Touretzky, H. S. Wan, and A. D. Redish. Neural representation of space in rats and robots. In J. M. Zurada and R. J. Marks, editors, *Computational Intelligence: Imitating Life*. IEEE Press, 1994. Proceedings of the symposium held at the 1994 IEEE World Congress on Computational Intelligence.