

Development of Abstract Categories in Embodied Agents

Giuseppe Morlino^{1,2}, Andrea Sterbini², Stefano Nolfi¹

¹ Institute of Cognitive Sciences and Technologies (CNR-ISTC), Roma, Italy

² Department of Computer Science, University “la Sapienza”, Roma, Italy
pinobrino@gmail.com, sterbini@di.uniroma1.it, stefano.nolfi@istc.cnr.it

Abstract. In this paper we demonstrate how a neuro-robot situated in an environment containing parallelepiped objects that vary in shape, size, and orientation can develop an ability to recognize and label the category of the objects and generalize to new objects. The analysis of the dynamical system constituted by the robot and the environment in interaction allowed us to understand how adapted agents solve the categorization problem at the level of the detailed mechanisms and at the level of the general strategy.

Keywords: Categorization, dynamical systems, evolutionary robotics.

1 Introduction

In this paper we demonstrate how a simulated neuro-robot situated in an environment containing parallelepiped objects that vary in shape, size, and orientation can develop through an evolutionary method [1] an ability to recognize and label the category of the objects (i.e. to discriminate whether the objects have a square or rectangular base). Since the sensors of the robots only provide information about a limited portion of the object, the categorization processes requires an ability to integrate sensory-motor information over time.

The goal of the paper is to study how a robot can associate sensory-motor values which vary continuously in state and time to abstract categories such as *square* or *rectangle* (a capacity that might represent a prerequisites for the development of several cognitive skills such as language [2]). Since the aim of the paper is not to investigate the role of active perception (i.e. how the possibility to co-determine the experienced stimuli through actions can be exploited to enable or to simplify the categorization process, c.f. [3-6]), the ability of the robot to explore the objects by traveling around them and to label their category have been evolved in two successive phases. The analysis of the coupled robot/environment system, also through the use of mathematical tools of dynamical system theory, allowed us to understand how the adapted agents solve the categorization problem both at the level of the detailed mechanisms and at the level of the general strategy. Moreover, the analysis conducted allowed us to elucidate the relation between the solution developed by the robots and the solution that might appear intuitive from the point of view of a human external

observer that consists in measuring and comparing the length of two adjacent sides of the object and that involves a form of relational categorization, c.f. [7].

2 The Experimental Scenario

The experiment involves a Khepera robot [8] situated in an arena that contains a square or a rectangular parallelepiped (Fig. 1). The robot is provided with eight infrared proximity sensors (which detect obstacles up to a distance of ~ 4 cm) and two motors which control the desired speed of the two corresponding wheels. The experiments have been carried out in simulation by accurately modelling the robot/environment interactions through a sampling technique [1].

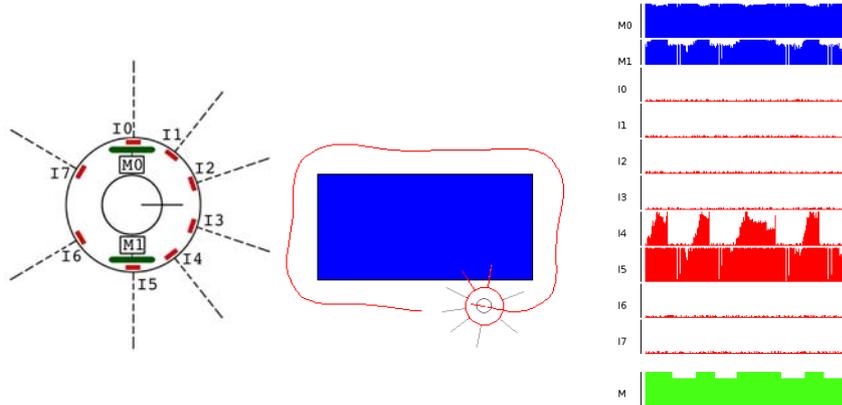


Fig. 1. Left: The position of the eight infrared sensors on the robot's body. Centre: The trajectory produced by an adapted robot that is traveling around a R object of 56x28 cm. Right: the activation of the left and right motors (M0-M1) controlling the two corresponding wheels and of the eight infrared sensors (I0-I7) during the behavior shown in the central picture. M is a neuron used in the second experiment that is set to 0.8 or 1.0 depending on whether M1 is below or above 0.9.

Each robot is evaluated for 40 trials during which it is allowed to interact with 20 square (S) and 20 rectangular (R) parallelepipeds of different sizes. The ratio between the length of the long and short sides in R objects is always $\frac{1}{2}$. The depth and the width of the objects vary for each trial within [20, 80] cm and are selected so to ensure that each side length occurs with the same probability in the S and R objects during the 40 trials. At the beginning of each trial, the robot is positioned at the center of the south side of the object, oriented towards west, and the state of the internal neurons of the robot (see below) is set to 0.0. Each trial lasts 1000 time-steps and each step lasts 100ms.

The robot's control system consists of a neural network composed by two modules: a motor module (Fig. 2, left) that regulates the speed of the two wheels, and a categorization module (Fig. 2, right) that determines the robot's categorization output (label).

The motor module is composed by eight infrared sensory neurons (I0-I7) directly connected to two motor neurons (M0-M1). The categorization module is composed by two neurons (M0-M1_(t-1)) that encode the state of the two corresponding motor neurons at time t-1), three internal neurons with recurrent connections (H0-H2), and one categorization output neuron (C).

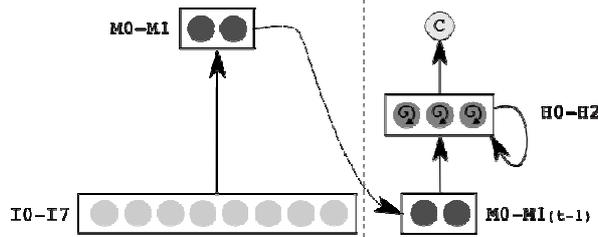


Fig. 2. The architecture of the robot's neural controller. The arrows between blocks indicate that all neurons of the second block receive connections from all neurons of the first block (or of the same block in the case of the three internal neurons).

The desired speed of the two wheels is set on the basis of the output of neurons M0-M1 normalized in the range [-10, 10]. The sensory neurons (I0-I7) are relay units which encode the state of the corresponding infrared sensors normalized in the range [0.0, 1.0]. The neurons (M0-M1_(t-1)) encode the state of the motor neurons at time t-1. The motor neurons (M0-M1) and the categorization unit (C) are updated according to the standard logistic function. The sensors, the neural controller, and the motors are updated every 100ms (the time step duration). The three internal neurons are leaky integrators which are updated on the basis of the following equations:

$$O_j = O_j^{(t-1)} \tau_j + \left(1 + e^{-A_j}\right)^{-1} (1 - \tau_j) \quad A_j = t_j + \sum_i w_{ij} O_i \quad (1)$$

With A_j being the activity of the j th neuron, t_j being the bias of the j th neuron, w_{ij} the weight of the incoming connections from the i th to the j th neuron, O_i the output of the i th neuron, $O_j^{(t-1)}$ the output of the j th neuron at the previous time step, τ_j the time constant of the j th neuron.

The characteristics of the robot's body and of the architecture of the robot's neural controller are fixed. The connection weights, the biases, and the time constant of the internal neurons are adapted through an incremental evolutionary method (Nolfi & Floreano, 2000) that includes two phases.

The initial population consists of 100 randomly generated genotypes which encode the free parameters of 100 corresponding individuals. Each parameter is coded with 8-bit and is normalized in the interval [-5.0, +5.0] for the biases and the synaptic weights, and in the interval [0.0, 1.0] for the time constants. Each subsequent population is obtained by selecting and retaining the best 20 individuals (the *elite*) of the previous population and by applying mutations (with 3% probability of flipping a bit) to 4 copies of each best individual.

During the first phase of the evolutionary process the free parameters of the motor neural module has been adapted for the ability to circumnavigate the object. More precisely, the fitness of the individuals has been calculated by computing the number or times the robot approaches a new corner of the object (i.e. every time the robot navigates from one corner to the next). This phase has been continued for 100

generations during which the robot develop an ability to circumnavigate objects of different size by displaying a wall following behavior (see Fig. 1, centre).

During the second phase the free parameters of the motor neural module are fixed on the basis of the parameters of the best individual obtained during the previous phase while the parameters of the categorization neural module have been evolved for the ability to label the category of the object. More precisely, the fitness of the individuals consists of the average absolute difference between the output of the categorization unit (C) and the desired value (i.e. 0.0 and 1.0 for S and R objects, respectively) calculated in each time step during the second half of each trial. The second phase is continued for 2000 generations and replicated 20 times starting from different randomly initialized genotypes.

2 Results

The analysis of the fitness at the end of the evolutionary process indicates that the evolved robots display close to optimal performance (more than 95% of correct categorizations) in 3 out of 20 replications, and good but sub-optimal performance in the other replications of the experiment. Moreover, we observed that the best adapted individuals display a remarkable ability to generalize their skill (within limits) to objects that differ, respect to those experienced during the adaptive process, either in size and/or in the ratio between their shorter and the longer sides.

Fig. 3 (left), shows the results of a test in which the best adapted robot is evaluated for 10,000 trials in a test condition in which we systematically varied the length of the north/south and east/west side within [10, 200] cm. As shown in the figure, in fact, the robot categorizes correctly S and R objects in the range [20, 100] cm.

The analysis also demonstrates that the adapted robot displays the two constituting properties of categorical perception: *labelling*, i.e. the capacity to partition stimuli varying in a continuous manner into well distinct classes, and *discrimination*, i.e. the tendency to better distinguishing between classes than within classes [5, 9-10]. More specifically, the presence of sharp boundaries between the two categories demonstrates that the robot partitions objects varying in a continuous manner into two well differentiated categories. As can be seen in the figure, objects are partitioned between the two classes on the basis of whether the ratio between their shorter and longer sides is higher or lower than ~ 0.7 in a way that is substantially independent from the size of the objects within the [20, 100] cm range. The sharp transition between the two categories shows that the output of the categorization unit (C) varies more for objects of different categories than for objects of the same category. These labelling and discrimination properties show that the robot categorizes objects (with side/ratios that differ from those experienced during the adaptive process) on the basis of the similarity between their sides/ratio and the sides/ratio of the S and R objects experienced during the adaptive process. In other words the analysis shows that the robot generalizes its skill also for objects with different sides/ratio. Moreover, the symmetry of the figure with respect to L1 and L2, that represent the length of the north/south and east/west sides, demonstrates that the robot categorizes rectangular

objects correctly independently from whether they are oriented vertically or horizontally.

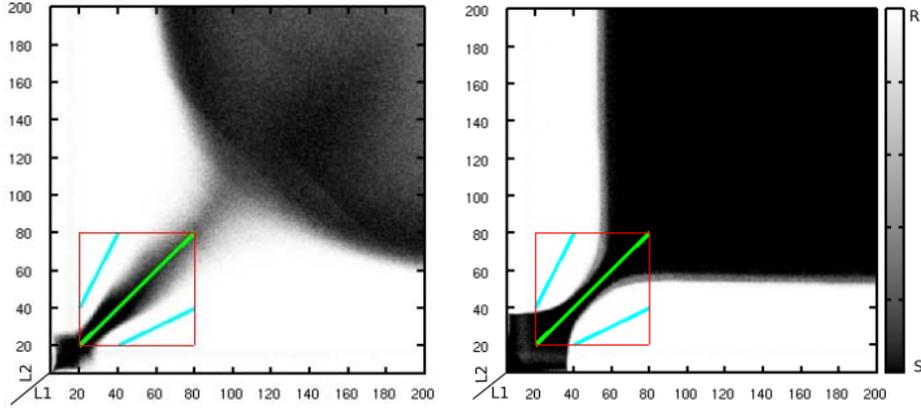


Fig. 3. Categorization outputs produced in interaction with objects with sides varying within [10, 200] cm. The L1 and L2 axis represents the length of the east/west and north/south sides of the objects, respectively. The greyscale colour represents the value of the categorization unit at the last time step averaged over 10,000 trials (0.0 = black, 1.0 = white). The central and the other two shorter lines inside the bottom-left square represent the size variation of the S and R objects experienced during adaptive process. Left: Results for the best individual of the basic experiment. Right: Results for the best individual of the simplified experiment (see section 2.1).

2.1 Dynamical analysis of the coupled robot-environmental system

To understand how the evolved robots categorize the two classes of objects we analysed the coupled dynamical system constituted by the robot and the environment. To overcome the problems due to the high dimensionality of the system we decided to analyse a slightly simplified version of the model. More specifically, in consideration of the fact that the two input neurons of the categorization module ($M0-M1_{(t-1)}$) encode redundant information, we ran a second experiment in which the categorization module includes only one neuron (M) whose activation state is set to 0.8 or 1.0 when the activation of the motor neuron $M1$ at time $t-1$ is lower or greater than 0.9, respectively (see Fig. 1, right).

The obtained results indicate that the simplified model leads to qualitatively and quantitatively similar results (see Fig.3, right). Fig. 4 displays the dynamics of the three internal neurons of the categorization module of the best adapted individual, produced when the state of the M neuron is fixed to 0.8 or 1.0 (left and right figures, respectively). Fig. 5 displays the typical trajectories of the same three internal neurons produced by the coupled dynamical system constituted by the robot and the environment in interaction.

One first thing to notice is that the state space includes four transient attractor points of which two (ACR/ASR) are located on the top and two (ACQ/ASQ) are located on the bottom of the state space. The attractors are transient since they manifest themselves when the M unit assumes one of the two possible values. More specifically, ACR and ACQ manifest themselves when the robot is negotiating a corner (i.e. when $M=0.8$) and ASR and ASQ manifest themselves when the robot is

travelling along a side of the object (i.e. when $M=1.0$). The ACR and ASR attractors are located nearby in the top part of the state space while the ACQ and ASQ attractors are well separated in the bottom part of the state space.

Secondly the top and the bottom part of the state space (with the exception of the small corner area near P_0) trigger the R and S categorization answers, respectively.

The third thing to notice is that the basin of attractions of ACR, ACQ and ASQ are confined in their relative areas (i.e. in the top part of the state space for ACR and in the bottom part for ACQ and ASQ). The basin of attraction of ASR, instead, extends from the bottom to the top part of the state space and can thus bring the internal state of the robot from the bottom to the top part of the state space.

The fact that at the beginning of the trial the internal state starts from P_0 and then move quickly toward the ASQ and ACQ transient attractors implies that, after few time steps, the robot starts to produce S as a default categorization answer. Moreover it implies that for squared objects the state of the internal neurons remains in the bottom part of the state space (from which an S categorization answer is produced) while for rectangular objects, at a certain point, the state moves along the ASR basin of attraction from the bottom to the top part of the state space (from which an R categorization answer is produced). The fact that the ASQ and ACQ basins of attraction are confined on the bottom part of the state space (i.e. the fact that they cannot bring the state of the internal neurons from the top to the bottom part of the state space) implies that the R categorization answer is irreversible.

Overall, these considerations imply that, to understand how categorization occurs we should understand the conditions that determine whether the state of the internal neurons remains in the bottom part of the state space or it moves and then remains on the top part of the state space. In other words the conditions that determine whether the robot keeps producing the default answer (S) or it starts producing the alternative answer (R).

When the robot travels along a side of the object, the state space is dominated by the ASQ and ASR attractor points. Since the state of the internal neurons is initially set to P_0 , however, during the initial phase of the trial the trajectory of the internal state is affected only by the basin of attraction of ASQ, that is located on the bottom part of the state space. As soon as the robot negotiates a corner of the object, the previous attractors are replaced by ACQ and ACR. Since the bottom part of the state space is dominated by the basin of attraction of ACQ, the state of the internal neurons then starts moving toward the ACQ attractor point. The periodic alternation of the two transient attractors thus leads to a stable or quasi-stable limit cycle, in which the state of the internal neurons move toward the ASQ and ACQ attractors located on two opposite sides of the bottom part of the state space. During the exhibition of this limit cycle the state of the internal neurons never fully reaches the two attractor points due to the limited time duration of each attractor, the inertial nature of the internal neurons, and the fact that the internal state moves more quickly toward the latter than toward the former attractor. Whether or not the basin of attraction of ASR succeeds in breaking this limit cycle dynamics and in bringing the internal state in the top part of the state space depends from the following factors.

The first factor concerns the fact that the position and the extension of the limit cycles that emerge from the robot/environmental interaction vary along the ACQ and

ASQ dimension depending on the length of the last sides of the object negotiated by the robot.

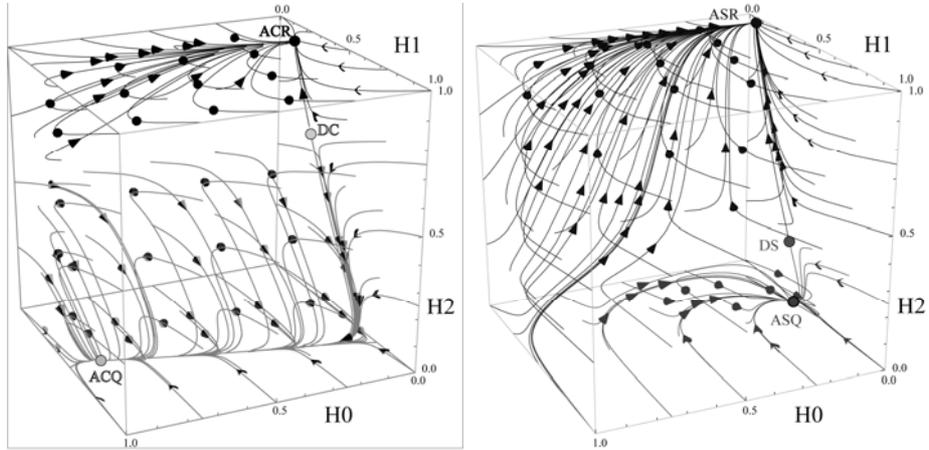


Fig. 4. Flow and phase portrait of the categorization module of the best adapted individual, produced when the state of M is fixed to 0.8 and 1.0 (left and right pictures respectively) that correspond to the states experienced by the robot when it travels along a corner or a side of the object, respectively. The three axes represent the state of the three internal neurons. The letter A stands for “attractor” (i.e. fixed point attractor), D for “saddle” (i.e. repellor), C for “corner”, S for “side”, Q for “cube/square” and R for “rectangle”; so for example ASR indicates the fixed attractor point that manifests itself when $M=1.0$ (when the robot is travelling along a side of the object) and that triggers a R answer.

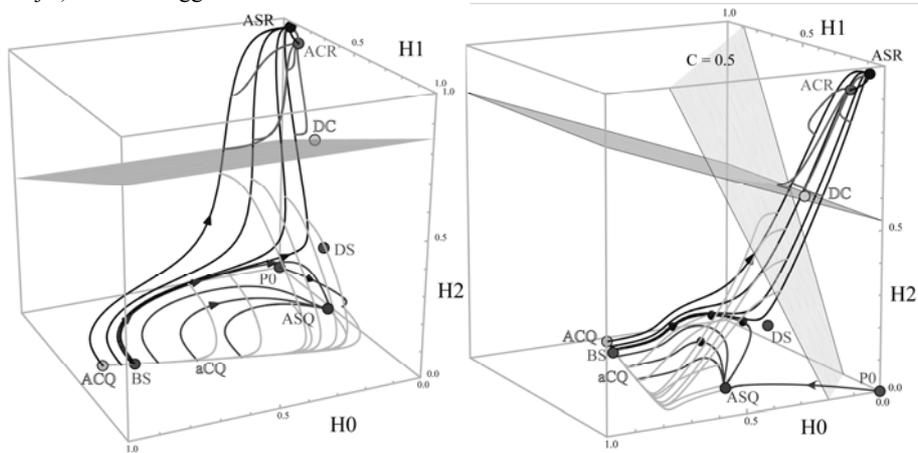


Fig. 5. Sample trajectories produced by the coupled dynamical system constituted by the robot and the environment in interaction. The left and right pictures display the same 3D structure from two different points of view. The three axes represent the state of the three internal neurons. The darker and lighter lines indicate the trajectory produced when the M neuron assumes a value of 0.8 or 1.0, respectively. P_0 indicates the initial state of the internal neurons. The decision plane that intersects the DC saddle point indicates the border between the basins of attractions of the ACR and ACQ attractor points. In the right picture, the categorization plan

indicated with $C=0.5$ separates the two areas of the state space that trigger an S or R categorization answers, respectively.

The extent to which the internal state approaches the ACQ attractor point primarily depends from the distance between the state of the internal neurons and the attractor at the beginning of the negotiation of the corner, which in turn is inversely proportional to the time duration of ASQ attractor that manifested itself while the robot was travelling along the previous side. We say “primarily” since the extension of the limit cycle toward the ACQ attractor also depends from the length of sides negotiated before the last one, thanks to the same effects described above for the last side (although the impact of previous stimuli tends to become progressively less important over time). The length of the last side negotiated by the robot (and, secondarily, the length of the sides negotiated before the last one) thus determines whether, while moving toward the ACQ attractor point, the internal state overcomes the BS point so to enter (after the negotiation of the corner) into the basin of attraction of the ASR attractor point located on the top part of the state space. The extent to which the internal state approaches the ACQ attractor also depends from the amount of time in which the attractor manifests itself that, however, is approximately the same for all corners (independently of whether they belong to square or rectangular objects).

The second factor that determines whether the state of the internal neurons crosses the decision plane that intersect the DC saddle point and enters (and then remains) into the top part of the state space or not, depends from the time duration of ASR (i.e. from the length of the current side) and from the extension of the limit cycle toward ACQ (that is inversely proportional to the length of the previous side, primarily, and of the sides before the previous, secondarily).

This type of analysis reveals also why the robot’s generalization ability is limited within the range described earlier. Additional information on this point as well as additional explanatory material is available from <http://laral.istc.cnr.it/esm/abstract-categorization/>.

3 Conclusions

In this paper we demonstrated how a simulated neuro-robot can develop an ability to associate sensory-motor stimuli which vary continuously in state and time to abstract categories. The analysis of the coupled dynamical system constituted by the robot and the environment in interaction demonstrates that the problem is solved by exploiting dynamical processes occurring at different time scales and the fact that the stimuli experienced by the robot can act as parameters that lead to sharp transitions in the robots’ internal dynamic. More specifically, the slow dynamic that originates from the inertial nature of the internal neurons allows the robot to detect and to keep trace in its internal state of the duration of previously experienced events (e.g. the time duration of the action produced by the robot along a side of the object). On the other hand, the sudden alternation of different type of stimuli lasting for certain time durations allows the robot to suddenly rearrange its internal dynamic in crucial phases (e.g. to perform an implicit comparison between the length of current and previous events). Moreover, it

is exploited to produce sharp transitions in the robot's internal dynamic (e.g. to channel the state of the internal neurons toward different areas of the state space associated to different categories).

At a more general level of description, the solution developed by the evolved robots demonstrates how the problem admits different solutions, including solutions that are more parsimonious and robust with respect to those that can be identified by an external observer. In particular, the problem faced by the robot is solved without fully partitioning the quantity to be compared (with particular reference to the length of the previous side) and by exploiting all available regularities (e.g. the overall size of the object and the fact that, in the domain of the experiment, very large and small objects always belong to the square category).

Acknowledgments. This research work was supported by the ITALK project (EU, ICT, Cognitive Systems and Robotics Integrating Project, grant n° 214668). The dynamical analysis as been carried on with Dynamica© [11]. The authors would also like to acknowledge Vito Trianni for his valuable insights on the analyses.

References

1. Nolfi, S., Floreano, D.: *Evolutionary Robotics: The Biology, Intelligence, and Technology of Self-Organizing Machines*. MIT Press/Bradford Books, Cambridge, MA (2000)
2. Harnad, S.: To Cognize is to Categorize: Cognition is Categorization. In: Cohen, H., Lefebvre, C. (eds.) *Handbook of Categorization in Cognitive Science*, Elsevier (2005)
3. Scheier, C., Pfeifer, R., Kuniyoshi, Y.: Embedded neural networks: exploiting constraints. *Neural Networks*, 11 (7-8), 1551–1596 (1998)
4. Nolfi, S., Marocco, D.: Active perception: A sensorimotor account of object categorisation. In: Hallam, B., Floreano, D., Hallam J., Hayes G., Meyer, J.-A. (eds.) *Proceedings of the 7th International Conference on Simulation of Adaptive Behavior*, Cambridge, MA, MIT Press, Cambridge, pp. 266–271 (2002)
5. Beer, R.D.: The dynamics of active categorical perception in an evolved model agent. *Adaptive Behavior*, 11(4), 209–243 (2003)
6. Nolfi, S.: Categories formation in self-organizing embodied agents. In: Cohen, H., Lefebvre, C. (eds.) *Handbook of Categorization in Cognitive Science*, Elsevier, pp. 869–889 (2005)
7. Williams P., Beer R., Gasser M.: An embodied dynamical approach to relational categorization. In: Love B., McRae K., Sloutsky, V. (eds.) *Proceedings of the 30th Annual Conference of the Cognitive Science Society*, Cognitive Science Society, Inc., pp. 223–228 (2008)
8. Mondada, F., Franzi, E., Ienne P.: Mobile Robot miniaturisation: A tool for investigation in control algorithms. In: *Proceedings of the Third International Symposium on Experimental Robotics*, Kyoto, Japan (1993)
9. Studdert-Kennedy, M., Liberman, A. M., Harris, K. S., Cooper, F. S.: Motor theory of speech perception: A reply to Lane's critical review. *Psychological Review*, 77, 234–249 (1970)
10. Ehret, G.: Categorical perception of sound signals: Facts and hypotheses from animal studies. In: Harnad, S. (ed.) *Categorical perception: The groundwork of cognition*. Cambridge, UK: Cambridge University Press, pp.301–331 (1987)
11. Beer, R.D.: *Dynamica Tutorial*. Version 1.0. First public release. (2008)