
**Neural Systems and Artificial Life Group,
Institute of Psychology,
National Research Council, Rome**

Emergence of Functional Modularity in Robots

Raffaele Calabretta, Stefano Nolfi, Domenico Parisi and Günter P. Wagner

In: R. Pfeifer, B. Blumberg, J.-A. Meyer and S.W. Wilson (eds.), *From Animals to Animats 5*, pp. 497-504, MIT Press (1998).

Department of Neural Systems and Artificial Life
Institute of Psychology, Italian National Research Council
V.le Marx, 15 00137 Rome - Italy
Phone: +39-06-86090233, Fax: +39-06-824737
E-mail: rcalabretta@ip.rm.cnr.it, parisi@ip.rm.cnr.it
<http://gral.ip.rm.cnr.it/rcalabretta>

Emergence of Functional Modularity in Robots

Raffaele Calabretta^{1,2}, Stefano Nolfi², Domenico Parisi² and Günter P. Wagner¹

¹ Department of Ecology and Evolutionary Biology, Yale University
New Haven, CT 06511, U.S.A.
e-mail: raffaele@peaplant.biology.yale.edu

² Department of Neural Systems and Artificial Life
Institute of Psychology, National Research Council
00137 Rome, Italy

Abstract

The origin and structural and functional significance of modular design in organisms represent an important issue debated in many different disciplines. To be eventually successful in clarifying the evolutionary mechanisms underpinning the emergence of modular design in complex organisms, one should be able to cover all different levels of biological hierarchy. Specifically, one should be able to investigate modularity at the behavioral level - the level on which natural selection operates - and understand how this level is related to the genetic level - the level at which natural selection works through mutation and recombination. We describe a simulation of the evolution of a population of robots that must execute a complex behavioral task to reproduce. During evolution modular neural networks, which control the robots' behavior, emerge as a result of genetic duplications. Simulation results show that the stepwise addition of structural units, in this case genetic and neural 'modules', can lead to a matching between specific behaviors and their structural representation, i.e., to functional modularity.

1. Introduction

The origin and structural and functional significance of modular design in organisms represent an important issue debated in many different disciplines such as evolutionary and developmental biology (e.g., Wagner, 1996; Raff, 1996; Wagner, 1995; Fontana and Buss, 1994; Needham, 1993; Bonner, 1988; Gould, 1977), the neurosciences, and cognitive science (e.g., Bates, 1994; Karmiloff-Smith, 1992; Fodor, 1983; Chomsky, 1957).

Many researchers tend to assume that human cognitive processes are accomplished by means of specialized modules (see e.g., Moscovitch and Umiltà, 1990) and, in fact, the modularity of mind is one of the most fundamental assumptions of cognitivism (Fodor, 1983). The meaning of the word «module», however, is rather problematic since the word itself "is used in markedly different ways by

neuroscientists and behavioral scientists, leading to considerable confusion and misunderstanding in interdisciplinary discussions" (Bates, 1994). This is just an example of a more general problem, that is, the existence of critical differences in the vocabularies and concepts used by the sciences which study organisms at different levels of organization: the biological sciences, the neurosciences, psychology, ecology, and the social sciences. The lack of a common conceptual apparatus represents a problem in that it makes more difficult to study the emergence of complex adaptations such as modularity, whose explanation may require the simultaneous analysis of phenomena and processes at many different levels (molecular, genetic, neural, behavioral, and population level).

In evolutionary biology, the concept of modularity is used to capture the fact that the bodies of higher organisms appear to be composed of semi-autonomous units (Raff, 1996; Wagner, 1996; Bonner, 1988). This observation raises several problems (Wagner, in preparation). First of all, the problem of identifying these 'natural kinds' (Quine, 1969; see also Griffiths, 1997; Wagner, 1996; Wagner, 1995), that is, the problem of 'decomposing' the organism into the appropriate units of phenotypic adaptation. (Notice that a similar problem exists in behavior-based robotics, where one may want to break down the desired complex behavior into an appropriate set of behavioral modules in order to design an integrated robot control system; see Nolfi, 1997a; Nolfi, 1997b; see also below). It has been argued that modularity is a prerequisite for the adaptation of complex organisms: modularity would allow the adaptation of different functions with little or no interference with other functions (Bonner, 1988). However, there is little understanding of how modularity originates, works, and remains incorporated in the genome.

In the field of Artificial Life, some researchers have tried to exploit modular design for improving the performance of various artificial systems such as artificial neural networks, evolutionary algorithms, and robots (see e.g., Nolfi, 1997a; Cho and Shimohara, 1997; Gruau, 1994). Conversely, evolutionary biologists are starting to appreciate the insights that could come from using these simulation models for the study of complex adaptations (see e.g., Levin *et al.*, 1997; Wagner and Altenberg, 1996; Frank, 1996; Toquenaga and

Wade, 1996; Wagner, 1995). The cross-fertilization between Artificial Life and Real Life could take place since Artificial Life partially shares the theoretical apparatus and vocabulary of evolutionary biology and, therefore, could supply it with additional methodological tools.

More specifically, models based on neural networks and genetic algorithms allow us to reproduce in a computer both the organisms and the environment in which they live, behave and reproduce. An organism can be simulated as having a body with a specific size, external shape, sensory and motor organs, etc., and an internal structure made up of a genotype, the nervous system, and other organs. Artificial neural networks that model organisms can be analyzed at the genetic level, at the level of the mapping from genotype to phenotype (development), at the neural and behavioral level, at the level of the effects of the network's output on the environment, at the level of the reproductive chances of each individual (fitness), at the level of populations of individuals and of entire ecosystems. Examining organisms at various levels could be crucial for understanding their behavior, because often an explanation of what happens at one level can be found at another level (see, for example, Miglino *et al.*, 1996; Calabretta *et al.*, 1996). As a consequence, one could hypothesize that the evolution of modularity results from the interaction among processes at different levels.

To evolve a neural controller for a mobile robot, Nolfi (1997a) used a modular neural network architecture that clearly outperformed other architectures in performing a task of garbage collecting (see below). It was denoted as an "emergent modular architecture" because it allows the required behavior to be broken down into sub-components corresponding to different neural modules as a result of an evolutionary adaptive process. Interestingly, Nolfi pointed out that evolved modules were correlated with a 'proximal' description of behavior, i.e., a description from the point of view of the robot's sensory-motor system that accounts for how the agent itself reacts in different sensory situations, rather than with a 'distal' description of behavior, i.e. a high level description in which terms such as 'approach' or 'avoid' are used to describe, from the observer's point of view, an entire sequence of sensory-motor steps (Sharkey and Heemskerk, in press). In other words, differently than in biology where distal functions are mapped to developmental modules, in evolved robots there was no correspondence between distal functional tasks and their structural representation.

It is important to stress that Nolfi was mainly concerned with exploiting modular design for synthesizing control systems for physical robots able to produce complex behavior. As a consequence, in his simulations the architecture of neural networks was pre-designed as modular right from the beginning and it remained fixed throughout the evolutionary process (*hardwired modularity*).

To investigate the issue of how modularity can emerge in nature, we present a modification of Nolfi's model (Nolfi, 1997a) in which gene duplication is also included as part of the evolutionary process and, therefore, modular neural




Figure 1 The Khepera robot.

networks can evolve starting from a population of non-modular ones as a result of gene duplication (*constructed* or, more specifically, *duplication-based modularity*). Our preliminary simulation results show that duplication-based modular architecture outperforms non-modular architecture, which represents the starting architecture in our simulations. Moreover, an interaction between mutation and duplication rate emerges from our results. Interestingly enough, duplicated modules, which are identical when duplication first takes place, begin to differentiate across generations as a result of genetic mutations and, eventually, a match between 'distal' modules and their genetic representation tends to emerge.

2. The model

We ran a set of simulations in which neural networks (Rumelhart and McClelland, 1986) are evolutionarily trained to control a mobile robot (see Figure 1) designed to keep an arena clear by picking up trash objects and releasing them outside the arena. The robot has to look for 'garbage', somehow grasp it, and take it out of the arena.

This task can be broken down into several sub-tasks: (a) to explore the environment, avoiding the walls; (b) to recognize a target object and to place oneself in a position such that the object can be grasped; (c) to pick up the target object; (d) to move toward the walls while avoiding other target objects; (e) to recognize a wall and to place oneself in a position with respect to the wall that allows the object to be dropped out of the arena; (g) to release the object. Moreover, these sub-tasks can be broken down into smaller components. It should be noticed that the task chosen is particularly well suited for studying the role of modularity because it involves different basic behaviors that may be implemented in different neural modules. The task requires a controller able to produce very different motor responses

for similar sensory states. As an example consider the robot in front of a target object. The robot should avoid or approach the object as a function of the presence or absence of a target on the gripper. If there is already a target on the gripper, the target object should be ignored. If there is no target on the gripper, the target object should be picked up. The only difference between the two situations is the state of one sensor out of 7. A modular neural network that can use different neural modules in different environmental situations should have an advantage with respect to a uniformly connected network in learning to produce very different motor responses for very similar sensory patterns.

The organism is a miniature mobile robot called Khepera, developed at E.P.F.L. in Lausanne (Mondada *et al.*, 1994; see Figure 1). The robot is supported by two wheels that allow it to move in various directions by regulating the speed of each wheel. In addition, the robot is provided with a gripper module with two degrees of freedom. The two arms of the gripper can move in parallel through any angle from vertical to horizontal while the gripper can assume only the open or closed position. The robot is also provided with eight infrared proximity sensors (six sensors are positioned on the front of the robot and two on the back) and an optical barrier sensor on the gripper capable of detecting the presence of an object between the two arms of the gripper. The infrared sensors allow the robot to detect obstacles to a distance of about 4 cm.

The environment is a rectangular arena 60x35 cm surrounded by walls containing 5 target objects. The walls are 3 cm in height and target objects are cylindrical boxes with a diameter of 2.3 cm and a height of 3 cm. The targets are positioned randomly inside the arena. To speed-up the evolutionary process a simulator was used (see Nolfi, 1997b).

In this work we compare the results obtained with three different neural network architectures (see Figure 2). In all cases the robot has 7 sensor neurons and 4 motor neurons. The first 6 sensory neurons are used to encode the activation level of the corresponding 6 frontal sensors of Khepera (the two back sensors are ignored) and the seventh sensory neuron is used to encode the light sensor on the gripper. On the motor side the 4 neurons respectively codify for the speed of the left and right wheels and for the triggering of the 'object pick up' and 'object release' procedures.

The first architecture (a) is a simple feedforward network with 7 input units encoding the state of the 7 sensors and four output units encoding the state of the four effectors. The input units are directly connected to the output units through 28 connection weights (plus 4 biases). This architecture is not divided into modules.

The second architecture (b) is a modular one and it has been called emergent modular architecture (Nolfi, 1997a) because it allows the required behavior to be broken down into sub-components controlled by different neural modules, although it does not require the designer to do such a partition in advance. (Notice that in this paper the emergent architecture is referred to as *hardwired modular*




Figure 2 Architectures (a) and (b) are shown on the left and right side, respectively. Architecture (a) is used in the nonmodular population. Architecture (b) is used in the two modular populations (i.e., in both hardwired and duplication-based modular populations). However, in the hardwired modular population two modules compete to gain control of each of the four actuators. In the duplication-based modular population individuals of the initial generation have only one module for each motor, that is, they initially have architecture (a). A second competing module may be added in individuals of successive generations as a result of the duplication operator. Another difference is that in the hardwired modular population competing modules have different random weights from the beginning, while in the duplication-based modular population, when a second competing module is introduced, the two competing modules have identical weights.

architecture). There are two modules for each of the four outputs (two wheels, the object pick up procedure, and the object release procedure). In any particular input/output cycle only one of the two competing modules can control the corresponding output. Each module includes two output units: a motor output unit and a selector unit. The motor output unit determines the speed of the corresponding wheel or whether or not the two procedures are executed. The selector unit determines the probability that the module will control the corresponding output. In other words, which of the two competing modules determines the output depends on which of the two competing selector units is more activated. Both the motor output unit and the selector unit of each module receives 7 connections (plus one bias) from the 7 sensory neurons.

The third architecture (b) is also modular and is denoted as *duplication-based modular architecture* because, in this case, the modules are not hardwired in the architecture from the beginning of evolution but they can be added during the evolutionary process. Each module, as in the case of hardwired architecture, consists of two output units (one motor output unit and one selector unit) which receive connections from the 7 sensors. At the beginning of the evolutionary process there is only one module for each of the four outputs, i.e., always the same module controls the corresponding output. However, during reproduction, modules may be duplicated (see below). Duplicated modules, which are exactly the same when duplication takes place, can differentiate across generations because of genetic

mutations.

A genetic algorithm (Holland, 1992) was used to evolve the connection weights of our neural networks. Each connection weight or bias is encoded as a sequence of 8 bits in the genotype. We begin with 100 randomly generated genotypes each representing a network with the same architecture and a different set of random connection weights. Each individual is allowed to 'live' for 15 epochs, each epoch consisting of 200 input-output cycles or actions. At the beginning of each epoch the robot and the target objects are randomly positioned in the arena. An epoch is terminated either after 200 actions or after the first object had been correctly released. At the end of life, the best individuals are selected for reproduction. The 20 individuals that have accumulated the highest 'score' during their lives generate 5 copies each of their neural networks. These $20 \times 5 = 100$ new robots constitute the next generation. The process is repeated for 1000 generations.

Reproduction consists in generating copies of an individual's genotype encoding the network's connection weights (we are assuming non-sexual reproduction in haploid populations) with the addition of random changes to some of the bits of the genotype sequence (genetic mutations) and, in the case of the duplication-based modular architecture, the duplication of a random selected neural module. Individuals were scored for their ability to perform the complete sequence of correct behaviors, i.e., for their ability to release objects correctly outside the arena. However, in order to facilitate the emergence of this ability individuals were also scored (even if with a much lower reward) for their ability to pick up targets (see Calabretta *et al.*, 1998).

In the present preliminary model the maximum number of duplicated modules allowed in the case of the duplication-based modular architecture is one for each motor output and no module-deletion operator was used. As a result, the hardwired modular architecture, already described in Nolfi (1997a), is the more complex architecture that can possibly evolve starting from architecture (a). However, the addition of competing modules during the course of evolution (instead than right from its beginning) that are initially identical to their competing module (instead of being completely unrelated) may produce qualitatively different results in the case of the hardwired and duplication-based modular architecture, respectively.

3. Results and discussion

We present the results of several simulations in which we compare a simple feedforward neural network, the hardwired modular architecture, and the duplication-based modular architecture (see Figure 2). In all simulations a mutation rate of 1% was used (i.e., 2% of the bits of the genotype randomly selected were replaced with a new randomly selected value). For the duplication-based modular architecture we investigated the performance obtained with a duplication rate of 0.02%, 0.03% and 0.04% (i.e., 0.02%,

0.03% and 0.04% of the modules were duplicated per replication). We ran 10 simulations for each of the 3 different architectures described above. Each simulation started with populations of 100 networks with randomly assigned connection weights and lasted 1000 generations.

Both populations with modules reach a higher fitness level than a population with only the basic architecture and no modules (Calabretta *et al.*, 1998). However, the two populations with modules do not differ in terms of overall fitness except that fitness growth is slightly slower in the population with duplication-based modules (Calabretta *et al.*, 1998). In order to demonstrate that modularity plays a critical role, we varied the duplication rate in the population with duplication-based modules, with the result that both average and peak performance decreased linearly with a decreased duplication rate until the advantage of modular design was lost (Calabretta *et al.*, 1998).

Let us now examine the behavior of a typical evolved individual with hardwired modularity and a typical evolved individual with duplication-based modularity.

Nolfi (1997a) pointed out that in hardwired modular architecture there was no correspondence between evolved modules and 'distal' description of behavior; rather, module switching and interaction was correlated with low-level sensory-motor mapping. He explained this by pointing out that the neural modules and the selection mechanism are represented homogeneously and evolve at the same time, and that, therefore, solutions in which both components are kept as simple as possible tend to be selected. Because of gene duplication (see Weiss, 1990; Ohno, 1970), a different evolutionary dynamic is involved in our model. In fact, if we look at the behavior of a typical evolved individual of Nolfi's simulations and a typical evolved individual of our simulations, an interesting difference emerges.

In the former case, the evolved individual described in Figure 5 always uses a single module to control the left wheel, the pick-up procedure, and the release procedure (LM, PU, and RL) while it uses both neural modules only for the right wheel (RM). Nolfi (1997a) stressed that "those two modules competing for the control of the right motor are both used in all the phases that can be described as distal sub-behaviors: when the gripper is empty and the robot has to look for a target (i.e., when sensor LB is off); when the gripper is carrying a target and the robot has to look for a wall (i.e., when sensor LB is on); when the robot perceives something and has to disambiguate between walls and targets (i.e., when the W/T graph shows the upper or bottom line); when the robot does not perceive anything (i.e., when the W/T graph does not show any line); when the robot is approaching a target (i.e., when sensor LB is off and the perceived object is a target); when the robot is approaching a wall (i.e., when sensor LB is on and the perceived object is a wall); when the robot is avoiding a target (i.e., when sensor LB is on and the perceived object is a target); when the robot is avoiding a wall (i.e., when the sensor LB is off and the perceived object is a wall)." In other words, he did not find one of the two neural modules competing for the





Figure 5 A typical evolved robot with a hardwired modular architecture. The top part of the figure represents the behavior of the robot in its environment. Lines represent walls, empty and full circles represent the original and the final position of the target objects, respectively, the trace on the terrain represents the trajectory of the robot. The bottom part of the figure represents the type of object currently perceived, the state of the motor of the two wheels, and the state of the sensors across 500 cycles. The 'W/T' graph shows whether the robot is currently perceiving a wall (top line), a target (bottom line), or nothing (no line). The 'LM', 'RM', 'PU', and 'RL' graphs show the state of the left and right motors and of the pick-up and release procedures. For each motor, the top portion of the graph shows the activation state (after the arbitration between component modules has been performed by the selector units) and the bottom part shows which of the two competing neural modules has control (the thickness of the black line at the bottom indicates whether the first or the second module has control: thin line = module 1; thick line = module 2). The graphs 'I0' to 'I5' show the state of the 6 infrared sensors. Finally, the 'LB' graph shows the state of the light-barrier sensor. The activation state of input and output units is represented by the height with respect to the baseline. (In the case of output units the activation state of the output units of the module that currently has control is shown). (Figure redrawn from Nolfi, 1997a).

control of the right motor to be responsible for specific distal sub-behaviors. On the contrary each sub-behavior was the result of the contribution of both neural modules (see Figure 5).

In the case of duplication-based modular architecture, on the other hand, the evolved individual described in Figure 6 uses both competing neural modules to control the left motor, the right motor, the pick-up procedure and the release procedure (LM, RM, PU, and RL). In this individual - a very successful one (it is able, in fact, to accomplish the correct sequence of behaviors for all the 5 objects present in the environment) - neural modules or combinations of neural modules appear to be responsible specific distal behaviors. In other words, this evolved individual tends to use different modules in different environmental situations. More precisely, both in the case of the left motor and in the case of the right motor, module switching is significantly correlated to the presence (i.e., when sensor LB is on) or

absence (i.e., when sensor LB is off) of a target object on the gripper (see Figure 6). Furthermore, in the case of the pick-up procedure output unit (PU), module switching takes place every time the robot has to pick up an object, while in the case of the release procedure output unit (RL) it happens when the robot has to pick up and to release the object (see Figure 6).

4. Conclusions

Our results suggest that the stepwise addition of structural units, in this case genetic and neural 'modules', can lead to a matching between distal functional tasks and their structural representation, i.e., to functional modularity. A comparison with Nolfi's simulations (Nolfi, 1997a) indicates that such a match is functionally not necessary (performance level is similar in the two conditions; Calabretta *et al.*, 1998), but it results from the stepwise addition of structural units.




Figure 6 A typical evolved robot with a duplication-based modular architecture (see Figure 5).

A possible explanation of this effect is that the structural units have already adapted prior to the duplication event and most likely they have adapted to participate in multiple functions. After duplication the modules tend to specialize for one of these functions ultimately leading to a close match between functional tasks and structural modules. According to this interpretation modularity is a side-effect of this mode of evolution rather than dictated by the nature of the task to which the robots adapt. This is different from an earlier proposal (Wagner, 1996) in which the selection regime itself is defining the division of tasks. It is also different from Altenberg's constructional selection mechanism in which a structural unit increases in frequency if it has fewer pleiotropic effects (Altenberg, 1995). (Pleiotropy is «the influence of the same genes on different characters», Futuyma 1998, p. 429). In our model the duplicated units all have pleiotropic effects but they tend to reduce these effects by functional specialization after duplication. We conclude that our model suggests a third possible mode for the origin of modularity, besides constructional selection (Altenberg, 1995) and selection against pleiotropic effects (Wagner, 1996): modularity as a secondary effect of duplication. Functional modules originally emerge as duplicated modules which subsequently diverge and specialize for some specific function.

In the present work we have described some preliminary results in support of this model of evolution of modularity.

We are currently engaged in further statistical analyses that should substantiate our hypothesis.

Acknowledgments

The financial support of the Italian National Research Council (National Advisory Committee for Biological and Medical Sciences: one-year sabbatical fellowship to RC and Bilateral Project 1998-1999 to RC and GPW) is gratefully acknowledged. RC would also like to acknowledge the support of Aurelio Romeo, Jeffrey R. Powell, Valerio Sbordoni, Riccardo Galbiati, Matthew Schlesinger and Cesare Giordano, the assistance of Ginger Booth and Jan Taschner, and the useful discussions with the members of the GPW's lab at Yale University during weekly meetings.

References

- Altenberg, L. (1995). Genome growth and the evolution of the genotype-phenotype map. In Banzhaf, W. and Eeckman, F. H., editors, *Evolution and biocomputation. Computational models of evolution*. Springer Verlag, Berlino-Heidelberg.
- Bates, E. (1994). Modularity, domain specificity and the development of language. *Discussions in Neuroscience*, 10:136-149.

- Bonner, J. T. (1988). *The evolution of complexity*. Princeton University Press, Princeton.
- Calabretta, R., Galbiati, R., Nolfi, S. and Parisi, D. (1996). Two is better than one: a diploid genotype for neural networks. *Neural Processing Letters*, 4:149-155.
- Calabretta, R., Nolfi, S., Parisi, D. and Wagner, G. P. (1998). An artificial life model for investigating the evolution of modularity. To appear in Bar-Yam, Y., editor, *Proceedings of the First International Conference on Complex Systems*, Nashua, NH, September 21-26, 1997. Addison-Wesley, Reading.
- Cho, S-B. and Shimohara, K. (1997). Emergence of Structure and Function in Evolutionary Modular Neural Networks. In Husbands, P. and Harvey, I., editors, *Proceedings of Fourth European Conference on Artificial Life*. MIT Press, Cambridge.
- Chomsky, N. (1957). *Syntactic structures*. Mouton, The Hague.
- Fodor, J. (1983). *Modularity of mind*. MIT Press, Cambridge.
- Fontana, W. and Buss, L. (1994). The arrival of the fittest: toward a theory of biological organization. *Bull. Math. Biol.*, 56:1-64.
- Frank, S. (1996). The design of natural and artificial adaptive systems. In Rose, R. and Lauder, G. V., editors, *Adaptation*. Academic Press, S. Diego.
- Futuyma, D. J. (1998). *Evolutionary Biology*. Sinauer, Sunderland.
- Gould, S. J. (1977). *Ontogeny and Phylogeny*. Harvard University Press, Cambridge.
- Griffiths, P. E. (1997). What emotions really are. The problem of psychological categories. University of Chicago Press, Chicago.
- Gruau, F. (1994). Automatic definition of modular neural networks. *Adaptive Behavior*, 2:151-183.
- Holland, J. H. (1992). *Adaptation in natural and artificial systems: an introductory analysis with applications to biology, control, and artificial intelligence*. MIT Press, Cambridge.
- Karmiloff-Smith, A. (1992). *Beyond modularity: A developmental perspective on cognitive science*. MIT Press, Cambridge.
- Levin, S., Grenfell, B., Hastings, A. and Perelson, A. (1997). Mathematical and computational challenges in population biology and ecosystems science. *Science*, 275:334-343.
- Miglino, O., Nolfi, S. and Parisi, D. (1996). Discontinuity in evolution: how different levels of organization imply pre-adaptation. In Belew, R. and Mitchell, M., editors, *Adaptive Individuals in Evolving Populations*. Addison-Wesley, Reading.
- Mondada, F., Franzi, E. and Ienne, P. (1994). Mobile robot miniaturisation: a tool for investigation in control algorithms. In Yoshikawa, T. and Miyazaki, F. eds., *Experimental Robotics III*, Lecture Notes in Control and Information Sciences. Springer-Verlag, London.
- Moscovitch, M. and Umiltà, C. (1990). Modularity and neuropsychology: implications for the organization of attention and memory in normal and brain-damaged people. In Schwartz, M. F., editors, *Modular Deficits in Alzheimer-type dementia*. MIT Press, Cambridge.
- Needham, J. (1993). On the dissociability of the fundamental processes in ontogenesis. *Biol. Rev.*, 8:180-223.
- Nolfi, S. (1997a). Using emergent modularity to develop control systems for mobile robots. *Adaptive Behavior*, 5:343-363.
- Nolfi, S. (1997b). Evolving non-trivial behaviors on real robots: a garbage collecting robot. *Robotics and Autonomous Systems*, 22:187-198.
- Ohno, S. (1970). *Evolution by gene duplication*. Springer Verlag, New York.
- Quine, W. V. (1969). Natural kinds. In Quine, W. V. editor. *Ontological relativity and other essays*. Columbia University Press, New York.
- Raff, R. (1996). *The shape of life. Genes, development, and the evolution of animal form*. University of Chicago Press, Chicago.
- Rumelhart, D. and McClelland, J. (1986). *Parallel distributed processing: Explorations in the microstructure of cognition*. MIT Press, Cambridge.
- Sharkey, N. and Heemskerk, J. (in press). The neural mind and the robot. In Browne, A. editor. *Current Perspective in Neural Computing*. IOP Press.
- Toquenaga, Y. and Wade, M. (1996). Sewall Wright meets Artificial Life. *Trends in Ecology and Evolution* 11:478-482.
- Wagner, G. P. (1995). The biological role of homologues: a

building block hypothesis. *Jb. Geol. Palont.* 195:279-288.

Wagner, G. P. (1996). Homologues, natural kinds and the evolution of modularity. *American Zoologist* 36:36-43.

Wagner, G. P. (in preparation). Appendages, fins, limbs and hands: homology and the hierarchical organization of the body.

Wagner, G. and Altenberg, L. (1996). Complex adaptations and the evolution of evolvability. *Evolution* 50:967-976.

Weiss, K. (1990). Duplication with variation: metameric logic in evolution from genes to morphology. *Yearbook of Physical Anthropology*, 33:1-23.