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

**A case study of the evolution of modularity:
towards a bridge between evolutionary biology, artificial life, neuro- and
cognitive science**

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

Technical Report NSAL-97013

Dicembre 1997
(revised March 6, 1998)

In: C. Adami, R. Belew, H. Kitano, and C. Taylor (eds.), *Proceedings of the Sixth International Conference on Artificial Life*, pp. 275-284, Cambridge, MA: MIT Press (1998).

Department of Neural Systems and Artificial Life
Institute of Psychology, C.N.R., Viale Marx 15 - 00137 - Rome, Italy
voice: +39 6 86090233 fax: +39 6 824737
e-mail: rcalabretta@ip.rm.cnr.it
<http://gral.ip.rm.cnr.it/rcalabretta>

A case study of the evolution of modularity: towards a bridge between evolutionary biology, artificial life, neuro- and cognitive science

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

¹ Department of Ecology and Evolutionary Biology, Yale University
165, Prospect Street - New Haven, CT 06511, USA
e-mail: rcalabretta@ip.rm.cnr.it

² Department of Neural Systems and Artificial Life
Institute of Psychology, C.N.R.
Viale Marx, 15 - 00137 Rome, Italy

Abstract

The existence of modules is recognized at all levels of the biological hierarchy. In order to understand what modules are, why and how they emerge and how they change, it would be necessary to start a joint effort by researchers in different disciplines (evolutionary and developmental biology, comparative anatomy, physiology, neuro- and cognitive science). This is made difficult by disciplinary specialization. In this paper we claim that, because of the strong similarities in the intellectual agenda of artificial life and evolutionary biology and of their common grounding in Darwinian evolutionary theory, a close interaction between the two fields could easily take place. Moreover, by considering that artificial neural networks draw an inspiration from neuro- and cognitive science, an artificial life approach to the problem could theoretically enlarge the field of investigation. The present work is the first one in which an artificial life model based on neural networks and genetic algorithms is used to understand the mechanisms underlying the evolutionary origin of modularity. An interesting problem that we will address in this paper is whether modules that start as repeated elements because of genetic duplication can develop to become specialized modules. A linear regression statistical analysis performed on simulation data confirms this hypothesis and suggests a new mode for the evolution of modularity.

Introduction

Various disciplines concerned with the study of organisms and their behavior find it useful to refer to 'modules' as components that play identifiable roles in systems at various levels and tend to maintain their identity over time. Although nonmodularity may also play a part in biological structure and function, the existence of modules is recognized at all levels of the biological hierarchy. The 'modularity of mind' is a well-known assumption of symbol-manipulation models of cognition. The mind is seen as composed by a multiplicity of modules that are specialized for various behavioral capacities and areas of activity. Neuroscientists recognize in the brain various types of units above the cellular level: columns, areas, systems, etc. In fact the total architecture of the brain appears to be a mosaic of interacting components with structural and functional specialization. Geneticists subdivide the DNA chain into genes that code for proteins and control the genotype-to-phenotype mapping. Modules

are also recognized at levels lower and higher than the gene level. At a lower level, genes are composed of triplets (*codons*) of bases (adenine, thymine, cytosine, guanine and uracil), each of which codifies for a specific amino acid. At a higher level, each gene codifies for a specific protein. The sequence of amino acids for each protein, as it is codified exactly in DNA, contains all the information to determine the three-dimensional structure on which the function of that protein finally depends (see for instance Creighton 1993 and Calabretta, Nolfi, and Parisi 1995). As stressed by Doolittle and Bork (1993), proteins are often composed by a limited group of modular elements (*domains*) that have spread and multiplied during evolution in ways that are starting to be understood. At the phenotypic level evolutionary biologists recognize homologous and analogous phenotypic traits in organisms belonging to different species or higher taxa, and repeated components in individual organisms, such as vertebrae in mammals (see Futuyma 1998, p. 669).

Given the postulated existence of modules at all these levels and their importance for describing and explaining both structure and process at each level, it is critical to understand what modules are, why and how they emerge, how they change, etc. To achieve this understanding it appears to be crucial to be able to coordinate modules existing at different levels of the biological hierarchy and to understand how modules at one level are related to those at other levels. This is made difficult by disciplinary specialization. The sheer amount of detailed empirical data that must be taken into consideration at each level, the heterogeneity of theoretical vocabularies and empirical methods used to study phenomena at different levels, and the great complexity of the between-level mappings, make it very difficult to clarify the relationships among modules at different levels in real organisms.

One possibility, then, is to study these problems in artificial organisms. Artificial Life studies all kinds of biological phenomena as they occur in artificial organisms and it can help us overcome many of the difficulties encountered in trying to relate modules at different levels. First, artificial organisms are simpler than real organisms. Second, simulations of biological phenomena at different levels can adopt a unified theoretical framework to facilitate inter-level conceptual dialogue. Finally, the

computer is a very powerful research instrument that allows us to observe and manipulate complex phenomena and nonlinear interactions among large number of entities at each level and between levels.

In this paper we adopt an Artificial Life approach in the hope that this approach can shed some useful light on modules at different levels and how they are related to each other.

Previous work

Research in the field of neuro- and cognitive sciences tends to assume that human cognitive process are accomplished by means of specialized modules (see e.g., Moscovitch and Umiltà 1990, Fodor 1983; for a critique of Fodor's point of view see Karmiloff-Smith 1992). Cowey (1981) and Kaas (1989) ask why the brain has so many visual areas. Ballard (1986) suggests that a limitation on the number of neurons compels the brain to adopt a modular architecture. Stevens (1994) maintains that «the complexity of human brain arises not from the complexity of its basic processing elements (the cortical module), or the richness of connections between modules, but simply in the number of the modules present». (For some connectionist simulations of modularity, see Jacobs, Jordan, and Barto 1991 and Rueckl, Cave, and Kosslyn 1989).

Even if the recognition of the existence and importance of modularity has a long historical tradition, there is little understanding of how modularity has originated. Evolutionary biologists ask whether modularity is an inherent property of organisms and thus not the result of evolution or it is the result of selection shaping the genotype-phenotype mapping function (see for instance Wagner 1995). The evolutionary implications of modular organization for development have been described by John Bonner in his book on the evolution of complexity (Bonner 1988). Modularity would allow the adaptation of different functions with little or no interference with other functions. Several population genetic models have been suggested in order to explain the evolutionary origin of modular design (e.g., Wagner and Altenberg 1996; Wagner 1996; Altenberg 1995) but our current knowledge is insufficient to assess the plausibility of these models.

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. Gruau (1994) applies a genetic algorithm to the synthesis of neural networks using cellular encoding as a new technology. This technology «can automatically and dynamically decompose a problem into a hierarchy of sub-problems, and generate a neural network solution to the problem. The structure of this network is a hierarchy of sub-networks that reflect the structure of the problem.» Snoad and Bossomaier (1995) consider «how genetic algorithms (GAs) and artificial neural networks (ANNs) (connectionist learning models) complement each other and how combining them (i.e. evolving artificial neural networks with a genetic algorithm), may give insights into the evolution of structure and modularity in biological brains.» Cho and Shimohara

(1997) investigate «the emergence of structure and functionality of modular neural networks through evolution.» The model they present is applied to a visual categorization task with handwritten digits.

In order to evolve neural controllers for mobile robots, Nolfi (1997) describes a modular neural network architecture that clearly outperforms other architectures in performing a garbage-collecting task (see below). This architecture is called an 'emergent modular architecture' because although modules are available from the beginning it is evolution that decides whether to use them or not by breaking down the required behavior into sub-components corresponding to different neural modules. In the present work we use the same simulation scenario of Nolfi (1997) but we add the genetic operator of gene duplication in order to explore the relationship between the evolutionary emergence of modularity and the phenomenon of gene duplication.

To our knowledge, the present work is the first one in which an artificial life model based on neural networks (Rumelhart and McClelland 1986) and genetic algorithms (Holland 1992) is specifically used to understand the mechanisms underlying the evolutionary origin of modularity.

Duplication-based modules

In the present paper we are concerned with modules that play a role in the genotype-to-phenotype mapping. More specifically, we are interested in the evolution of modules at the genetic level that map into single functions at the behavioral level of the entire organism. Mappings from genes to higher functions can be modular or nonmodular (Wagner and Altenberg 1996). The mapping is modular when there are few pleiotropic effects among characters serving different functions, with pleiotropic effects existing mainly among characters which serve one and the same function (Figure 1, right). (Pleiotropy is «the influence of the same genes on different characters», Futuyma 1998). On the contrary, we have a nonmodular mapping when there are pleiotropic effects both among characters serving different functions and among characters serving a single function (Figure 1, left). Therefore, modules can be defined as a collection of characters at different levels that are all responsible mainly for a single function. Put simply, in the genes-to-behavior mapping a module can be defined as a collection of genes which produce a set of molecules which in turn are responsible in the regulation of the nervous system serving a given behavioral function. Notice how this definition of module is more constrained than others. Neuro-physiologists, for instance, in defining a module take into account the nervous system and the higher level of organization (behavior) which is the result of the activity of the nervous system. However, they do not usually take into consideration lower levels such as the molecular and genetic level. They do not ascertain that what they have identified as a neural module is the result of a collection of genes that mainly codify for that phenotypic character. If we take an evolutionary perspective, however, the genotype level plays a very important role because it is at




Figure 1. Examples of nonmodular and modular genotype-to-phenotypes mapping. Complexes of phenotypic characters {A, B, C} and {D, F, G, H} serve behavioral functions F1 and F2, respectively. The genetic representation is modular in the case to the right because some genes (i.e., {G1, G2, G3}) have primarily pleiotropic effects on the first set of characters (C1) supporting behavioral function F1 whereas other genes (i.e., {G4, G5, G6}) have primarily pleiotropic effects on the characters (C2) subserving function F2. The left case is nonmodular because there are about the same amount of pleiotropic effects on the characters subserving both functions. (Figure redrawn from Wagner and Altenberg 1996).

this level that novelties are produced through mutation, recombination, and selection.

Modules can be seen as specialized components and, therefore, different from each other, or they can be recognized as repeated identical elements. An interesting problem that we will address in this paper is how the two types of modules are related. In particular we will ask if modules that start as repeated elements because of genetic duplication can develop to become specialized modules.

Wagner and Altenberg (1996) stressed that «although modularity may sometimes be intrinsic to the mechanism of an organismal function, in many cases, especially development, modularity appears to be an evolved property.» A possible mechanism of morphological innovation is the differentiation of repeated elements (Müller and Wagner 1991; Ohno, 1970; Weiss 1990), for instance the differentiation of metameric segments at the origin of insects (see for instance Akam, Dawson, and Tear 1989). Various authors have stressed the role of genetic duplication for the emergence of evolutionary novelties, especially in complex organisms. Li (1983) claims that «gene duplication is probably the most important mechanism for generating new genes and new biochemical processes that have facilitated the evolution of complex organisms from primitives ones». Tautz (1992) argues that «redundancy of gene actions may [...] be a necessary requirement for the development and evolution of complex life forms» and in fact «redundancy seems to be widespread in genomes of higher organisms» (Nowak et al.

1997). In the neutral theory of molecular evolution (Kimura 1983), the duplication relaxes the selective constraints on one of the two copies allowing the accumulation of mutations leading to the emergence of a new function (Coissac, Maillier, and Netter 1997; see also Ohta 1989).

In the present work we present simulations of the evolution of populations of artificial organisms focusing on the evolutionarily emergence of functionally different modules at the neural-behavioral level from gene duplication.

A typical Artificial Life simulation addressing problems at the behavioral level involves a population of organisms living and reproducing in an environment. The behavior of each individual organism is controlled by a neural network that encodes the state of the local environment in its input units and some movement of the organism in its output units. Each individual has an inherited genetic code that specifies (some of) the properties of the individual's neural network and, therefore, of the individual's behavior. The individuals that inherit better neural networks tend to behave more efficiently and are more likely to leave offspring. The genetic code is inherited with random mutations and/or sexual recombination of parts of the genetic code of one parent and parts of that of the other parent. The resulting offspring are in many cases worse than their parents but, although infrequently, they can represent an improvement over their parents. The selective reproduction of the best individuals and the constant

addition of variability through mutations and/or sexual recombination make it possible to observe evolutionary change in the population at three levels: genetic, neural, and behavioral (Miglino, Nolfi, and Parisi 1996).

We compare two populations. In both populations neural modules start as reduplications in the genetic code and they evolve their connection weights during the evolutionary process. In one population the genetic code is hardwired from the beginning for coding for two distinct neural modules for each separate aspect of the network's output. In principle each of the two modules can control the same network's output. In the other population the emergence of distinct modules becomes an adaptive process in the sense that the genetic code includes a 'reduplication gene' that can be turned on at some point during the evolutionary process. An important difference between the two populations is that in the first population the two alternative neural modules controlling the same network's output both start from zero, i.e., from random connection weights, and they must evolve their connection weights in parallel to become specialized for different tasks, whereas in the second population a duplicated module starts with the weights already evolved for the first module and must then adapt these weights to differentiate and specialize with respect to the first module. We will call the first type of modules «hardwired» and the second type of modules «duplication-based».

The two populations are compared with respect to how much modules at the genetic level map into meaningful units at the behavioral level. More specifically we want to test the prediction that modular architectures that originate in genetic duplication tend to have modules corresponding to meaningful behavioral units more often than architectures with hardwired modules.

Let us explain what it is for a module to correspond to a meaningful behavioral unit. Imagine a population of organisms (robots) living in a walled environment that contains a certain number of objects. The task for these organisms is to grasp the objects with their 'arms' and to release the objects over the peripheral wall outside the environment. The entire behavioral sequence that allows the organisms to accomplish this task can be divided up into a hierarchy of meaningful units. At the highest level of the hierarchy the sequence can be divided into two units: grasping an object and releasing the object beyond the wall. At the next lower level, in order to grasp an object the organism must find the object and in order to do so it must discriminate the object from the peripheral wall, approach and reach the object. At the lowest level the organism must explore the environment until it perceives an object. Also releasing the object on the other side of the wall can be divided into subsegments: avoid and ignore the other objects (since only one object can be grasped by the organism's arms), reach the wall, open the arms to release the object beyond the wall. Each of these segments is a meaningful behavioral unit. Our question is whether neural modules specialize for these units in the sense that different modules are used when a particular behavioral unit must be executed. We believe that this may be so for modules that emerge from genetic duplication and represent evolutionary

specializations of already existing and functional modules whereas hardwired modules tend to be less clearly associated with meaningful behavioral segments.

Simulations

We ran a set of simulations in which two different populations of neural networks are trained to control a mobile robot 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 with its arms, and take it out of the arena.

The robot is a miniature mobile robot called Khepera, developed at E.P.F.L. in Lausanne (Mondada, Franzi, and Jenne 1993). 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 through any angle from vertical to horizontal while the gripper can assume only the open or closed position. The robot is also provided with six infrared proximity sensors positioned on the front of the robot 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 and containing 5 objects. The walls are 3 cm in height and the objects are cylinders with a diameter of 2.3 cm and a height of 3 cm. The 5 objects are positioned randomly inside the arena. To speed up the evolutionary process a simulator of the physical robot and environment was used (see Nolfi 1997).

The basic network architecture is identical in the two populations (see Figure 2). The architecture includes 7 input units directly connected to 4 output units, each with its associated bias, for a total of $(7 \times 4) + 4 = 32$ connections. Six of the 7 input units continuously encode the activation level of the 6 infrared sensors while the seventh input unit binarily encodes whether (1) or not (0) there is an object between the two arms of the gripper. Two of the 4 output units continuously encode the speed of Khepera's two wheels. The remaining 2 output units binarily encode whether (1) or not (0) each of two procedures are executed by the robot: one output unit encodes the procedure of picking up an object and the other unit the procedure of releasing the object.

The two populations differ in the type of modularity that enriches this architecture (see Figure 2). In one population the architecture of all individual organisms includes two modules for each of the 4 output units since the beginning of evolution. More specifically, the architecture has two copies for each of the 4 output units, with each copy receiving its own set of connections from the input units. Which of the two alternative output units actually controls the robot's behavior in each particular input/output cycle is decided in the following way. Each copy of an output unit has associated with it a special unit called a 'selector' unit that receives connections from all the input units and has its own bias. In each cycle the simulator ascertains which of

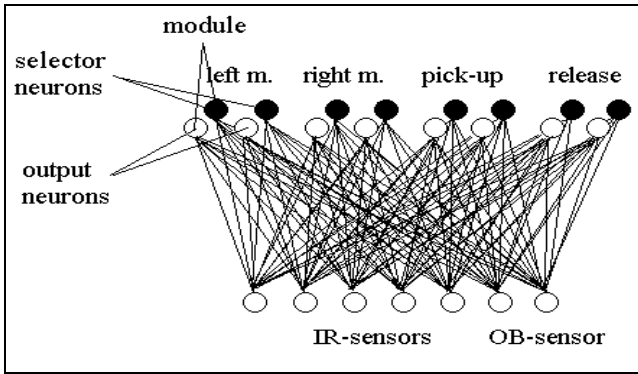


Figure 2. Modular neural network architecture of the two populations. The basic architecture is identical in the two populations. The two populations differ in the type of modularity which is added to this basic architecture. In one architecture two modules compete to gain control of each of the four actuators in all individuals since the beginning of evolution. In the second population the individuals of the initial generation have only one module for each motor. A second competing module may be added in individuals of later generations as a result of the duplication operator (see below). Another difference is that in the first population competing modules have different random weights at the beginning while in the second population when a second competing module is generated, the two competing modules have identical weights.

the two selector units is more activated and it uses the output unit corresponding to the more highly activated selector unit to determine the organism's behavior. One copy of each output unit, with its associated connections, plus its selector unit with its associated connections, constitute a module. For each output unit, therefore, there are two alternative modules that compete for controlling the organism's behavior and it is the input from the environment that ultimately decides which of the two alternative modules control the robot's behavior.

A genetic algorithm (Holland 1992) was used to evolve the connection weights of such neural networks. In the first population the genotype encodes the values for all the connection weights of the modular architecture. Since each module includes 7×2 connections plus 2 biases and there are 8 modules, the total number of connection weights encoded in the genotype is 128. Since each weight value is binarily encoded using 8 bits, the total genotype is a sequence of $128 \times 8 = 1024$ bits. The individuals of the first generation are assigned random values for these 1024 bits and then the evolutionary process progressively finds better and better genotypes on the basis of the selective reproduction of the best individuals and the addition of random mutations to inherited genotypes. Each generation includes 100 individuals. At the end of life the 20 best individuals are selected for reproduction and each of these individuals generates 5 offspring, that is, new individuals with the same genotype of their parent (reproduction is nonsexual). Genetic mutations consist in changing the value of about 10 bits in each genotype (1% mutation rate). The $20 \times 5 = 100$ new individuals constitute the second generation. The process is repeated for 1000 generations.

In the second population the genotypes of the initial

generation encode random values for the connection weights of the single modules of the basic architecture : $32 (7 \times 4 = 28 \text{ plus } 4 \text{ biases})$ connections. However, since each of the 4 output units has associated with a nonfunctional selector unit with its 7 connection weights, the total number of connection weights encoded in the genotypes of the initial generation is 64. Notice however that until the module is not duplicated this selector unit remains completely nonfunctional and its associated connection weights are subject to random drift only. The genotype of this second population has 4 additional 'duplication genes' each associated with one of the 4 output units. When one of these duplication genes is turned on by some mutation the gene duplicates its corresponding module assigning to the duplicated module the same weight values of the original module. The duplication genes cause a duplication with some probability that we have varied in various simulations (i.e., 0.04%, 0.03% and 0.02% of the modules were duplicate in different simulations). In the generation in which the duplication of some module occurs there is no possible change in behavior since both alternative modules have the same connection weights but subsequently random mutations acting on the module's connections weights (both on those leading to the output unit and those leading to the selector unit of the module) can progressively differentiate the two alternate modules. (As in the first population, we used a mutation rate of 1%).

In conclusion, we have two populations. One population has a fixed, hardwired modular architecture since the beginning of the evolutionary process. What we can determine with respect to this first population is, first, whether the evolved individuals do actually make use of the alternate modules as a function of the circumstances or they only use a single module for all environmental inputs, and second, in the case they use alternate modules, whether or not we can attribute a functional meaning to the modules, i.e., whether or not distinct modules control meaningful behavioral units. The other population starts with a nonmodular architecture but it is free to evolve a modular architecture if that turns out to be adaptive. In the present model modules can be evolutionarily added to neural architectures (with a limit of one module for each motor output) but they cannot be deleted. Hence, because of purely random reasons the individuals in this second population will tend to approximate the modular architecture of the first population, with two alternate modules for each output unit. However, the modules of the second population have a different origin than those of the first population. Not only are they evolved rather than hardwired but while the modules of the first population all start with random weights and therefore two alternate modules for the same output unit both evolve from zero (random connection weights), the alternate modules in the second population start with the same weights of the original modules (since they duplicate these modules) and therefore with weight values that are already adapted. What we want to determine is if the different origin and evolutionary history of modules that arise out of genetic duplication results in modules endowed with a greater amount of functional meaning at the behavioral level.

Results

Both populations with modules reach a higher fitness level than a population with only the basic architecture and no modules (cf. Nolfi 1997 and Calabretta et al. 1997). 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 (results not showed). 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 (see Calabretta et al. 1997).

We then examined the behavior of a typical evolved individual with hardwired modularity and a typical evolved individual with duplication-based modularity and found that an interesting difference emerged between the two individuals. While in the hardwired modular individual there was no correspondence between modules and meaningful behavioral units ('distal' description of behavior, according to Nolfi's definition), in the individual with duplication-based modularity neural modules or, better, combinations of neural modules turned out to be responsible for specific meaningful behavioral units (see Calabretta et al. 1997, figure 5 and figure 6).

In order to extend and reinforce this result we examined the best individual of the last generation in each of the 10 replications of the simulation for both populations and we compared the results concerning the statistical relationships between meaningful behavioral units and the use of the modules. Specifically, we considered as a meaningful behavioral unit the fact that the robot had or did not have a target object on the gripper. We tested the best individuals of the last generation in 10 different repetitions of the simulation for both populations. Each individual was allowed to live for 1 epoch consisting of 500 actions.

For each action we binarily recorded both the state of the modules (i.e., which of the two available modules for each motor output was active) and if the meaningful behavioral unit was being executed or not. For each repetition of the simulation we calculated the linear regression between meaningful behavioral unit as a categorical dependent variable and the state of modules as a categorical independent variable. As we already have said, we wanted to test the prediction that modular architectures that originate in genetic duplication tend to have modules corresponding to meaningful behavioral units more often than architectures with hardwired modules.

Table 1 shows the chi-square values for each repetition of the simulation both in the case of hardwired modularity and of duplication-based modularity. If we look at the frequency distribution of chi-square values, two distinct pictures emerge for the two models (see Figure 3). For the hardwired modularity model chi-square values are very low in 8 out of 10 replications of the simulation; more precisely, these values are less than 20 in 5 replications and less than 30 in 3 replications (see left graph of Figure 3 and also Table 1).

	Hardwired modularity's	Duplication-based modularity's
Seed	chi-square values	chi-square values
1	11.135	368.662
2	4.679	246.374
3	425.927	495.961
4	2.747	218.359
5	21.556	190.511
6	439.391	55.947
7	16.647	55.246
8	2.348	296.993
9	29.078	32.334
10	27.081	321.769

Table 1. Chi-square values for the single best individuals of the last generation in each repetition (initial random seed) of the simulation for hardwired modularity (left) and duplication-based modularity (right).

In other words, there is a very low correlation between the meaningful behavioral unit we have selected for examination and the use of specific modules in 8 out of 10 replications of the simulations (in 4 replications of the simulations the correlation is not significant at all). Modules do not appear to be specialized for the specific meaningful behavioral unit we have considered. Conversely, for the duplication-based modularity model chi-square values are very high in 9 of 10 replications of the simulation; more precisely, they are higher than 100 in 7 replications and higher than 50 in 2 replications (see right graph of Figure 3 and Table 1). In statistical parlance, the dependent variable (i.e., the meaningful behavioral unit) can be said to be a function of the independent variable (i.e., the state of modules), that is, there is a significant correlation between the considered meaningful behavioral unit and the usage of modules in all the 10 replications of the simulation. (Notice that the degrees of freedom and the significance values vary in different simulations depending on how many modules are functional in particular neural networks). This means that combinations of neural modules are specialized for the specific meaningful behavioral unit we have considered and that evolved individuals tend to use different modules in different environmental situations. In other words, the prediction that modular architectures originating in genetic duplication tend to have modules corresponding to meaningful behavioral units more often than architectures with hardwired modules appears to be confirmed by the present results.

Interpretation and Conclusions

The results presented above are suggestive of a new mode of evolution for modularity. Modularity may critically depend on the duplication and subsequent divergence of units that are already partially adapted to some functional task. This proposed mechanism is thus different from the combination of directional and stabilized selection on preexisting characters proposed in Wagner (1996) as well as from the 'constructional' selection for genes with lower degrees of pleiotropy proposed by Altenberg (1995).




Figure 3. Frequency distribution of chi-square values shown in the Table 1, both in the case of hardwired modularity (left) and of duplication-based modularity (right).

We suggest the following scenario to explain the results of our simulations. In our model, the evolution of functional specialization depends on the partial adaptation of the units prior to the duplication event. We tested this by simulating the addition of neural units with random connection weights. The results of these simulations show that this does not lead to the origin of functionally specialized modules (results not shown). We assume that prior to duplication the units serve more than one function. We further assume that these multiple functions lead to functional conflicts in the optimization of functional performance. A duplication of a multi-functional unit then releases these constraints. Consequently the duplicated units are free to specialize for one of the functions and a modular mapping between functions and neural modules emerges. We are currently undertaking simulations to test this hypothesis.

This interpretation of our simulation results is similar to one model of evolution by gene duplication which has been proposed by Hughes (1994). The standard model, going back to Ohno (1970), assumes that the gene has only one function prior to duplication but that after duplication one copy is free to explore new functional opportunities. It has been argued that this model is problematic in assuming that new functions can be acquired by random search, i.e., mutation and random drift. An alternative model proposed by Hughes (1994) assumes that prior to duplication the gene is serving multiple functions, and that the performance of these functions is not optimal because of conflicting adaptive demands. After gene duplication, the two copies are released from the conflicting functional demands and each gene copy specializes for one of the functions of the ancestral gene. This model is supported by the preponderance of evidence about the functional history of duplicated genes (Hughes, 1994).

If correct, this interpretation about the origin of functional modularity raises important questions about the relationship between evolutionary mechanisms and evolvability. As emphasized by Bonner (1988) and Wagner and Altenberg (1996), modular genetic architectures are superior in their ability to produce functionally improved mutations. But the question remains whether these genetic

architectures arise because of their impact on evolvability. There are a number of difficulties associated with the idea that evolvability arises as an adaptation to evolvability (for a recent discussion see Steward 1997). Our results further accentuate these problems, since the mechanism for the origin of modularity in our model does not derive from or is related to evolvability. Modularity appears to be a consequence of the evolution of functional specialization. Evolvability per se does not seem to be a factor in its origin. If this interpretation is correct, evolvability has to be seen as a secondary consequence of adaptation (effect selection) and not an adaptation to the evolvability of complex organisms.

Acknowledgments

The financial support of the Italian National Research Council (National Advisory Committee for Biological and Medical Sciences: one-year fellowship to RC and Bilateral Project 1998-1999 to RC and GPW) is gratefully acknowledged. RC would also like to acknowledge the assistantship of John W. Emerson of the Social Science Statistics Lab at Yale University for statistical analyses, and the useful discussions with the members of the GPW's lab at Yale University during weekly meetings and with the members of the Research Group on Artificial Life (GRAL) in Rome.

References

- Akam, M., Dawson, I., and Tear, G. 1989. Homeotic genes and the control of segment diversity. *Development* 104:123-133.
- Altenberg, L. 1995. Genome growth and the evolution of the genotype-phenotype map. In Banzhaf, W., and Eeckman, F. H. *Evolution and biocomputation. Computational models of evolution*. Berlino-Heidelberg: Springer Verlag.
- Ballard, D. H. 1986. Cortical connections and parallel processing: structure and function. *The Behavioral and Brain Sciences* 9:67-120.

- Bonner, J. T. 1988. *The evolution of complexity*. Princeton, New Jersey: Princeton University Press.
- Calabretta, R., Nolfi, S., Parisi, D., and Wagner, G. P. 1997. Evolutionary mechanisms for the origin of modular design in artificial neural networks, Technical Report, CCE-#51, Center of Computational Ecology, Yale University.
- Calabretta, R., Nolfi, S., and Parisi, D. 1995. An artificial life model for predicting the tertiary structure of unknown proteins that emulates the folding process. In Moran, F., Moreno, A., Merelo, J.J., and Chacon, P. eds. *Lecture notes in artificial intelligence: advances in artificial life.*, 862-875. Berlino-Heidelberg: Springer-Verlag.
- Cho, S-B., and Shimohara, K. 1997. Emergence of structure and function in evolutionary modular neural networks. In Husbands, P., and Harvey, I. eds. *Fourth European Conference on Artificial Life*. Cambridge, Mass.: MIT Press.
- Coissac, E., Maillier, E., and Netter, P. 1997. A comparative study of duplications in bacteria and eukaryotes: the importance of telomeres. *Molecular Biology and Evolution* 14:1062-1074.
- Cowey, A. 1981 Why are there so many visual areas? In Schmitt, F. O., Warden, F. G., Adelman, G., and Dennis, S. G. eds. *Models of the visual cortex*, 54-61. New York, New York: John Wiley and Sons.
- Creighton, T. E. 1993. *Proteins: structures and molecular properties*. New York, New York: W. H. Freeman and Company.
- Doolittle, R. F. and Bork, P. 1993. La modularità delle proteine nell'evoluzione. *Le Scienze* 304:58-64.
- Fodor, J. 1983. *Modularity of mind*. Cambridge, Mass.: MIT Press.
- Futuyma, D. J. 1998. *Evolutionary biology*. Sunderland, Mass.: Sinauer.
- 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*. Cambridge, Mass.: MIT Press.
- Hughes, A. L. 1994. The evolution of functionally novel proteins after gene duplication. *Proceedings of Royal Society. Series B* 256:119-124.
- Jacobs, R. A., Jordan, M. I., and Barto, A. G. 1991. Task decomposition through competition in a modular connectionist architecture: the what and where vision task. *Cognitive Science* 15:219-250.
- Kaas, J. H. 1989. Why does the brain have so many visual areas? *Journal of Cognitive Neuroscience* 1:121-135.
- Karmiloff-Smith, A. 1992. *Beyond modularity: a developmental perspective on cognitive science*. Cambridge, Mass: MIT Press.
- Kimura, M. 1983. *The neutral theory of molecular evolution*. Cambridge, UK: Cambridge University Press.
- Li, W-H. 1983. Evolution of duplicate genes and pseudogenes. In Masatoshi, M., and Koehn, R. K. eds. *Evolution of genes and proteins*. Sunderland, Mass.: Sinauer.
- 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. eds. *Adaptive Individuals in Evolving Populations*. Reading, Mass.: Addison-Wesley.
- Mondada, F., Franzi, E., and Ienne, P. 1993. 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. ed. *Modular Deficits in Alzheimer-type dementia*. Cambridge, Mass.: MIT Press.
- Müller, G. B., and Wagner, G. P. 1991. Novelty in evolution: restructuring the concept. *Annual Review of Ecology and Systematics* 22:229-256.
- Nolfi, S. 1997. Using emergent modularity to develop control systems for mobile robots. *Adaptive Behavior* 5:343-363.
- Nowak, M. A., Boerlijst, M. C., Cooke, J., and Maynard Smith, J. 1997. Evolution of genetic redundancy. *Nature* 388:167-171.
- Ohno, S. 1970. *Evolution by gene duplication*. New York, New York: Springer Verlag.
- Ohta, T. 1989. Role of gene duplication in evolution. *Genome* 31:304-310.
- Rueckl, J. G., Cave, K. R., and Kosslyn, S. M. 1989. Why are «what» and «where» processed by separate cortical visual systems? A computational investigation. *Journal of Cognitive Neuroscience* 1:171-186.
- Rumelhart, D., and McClelland, J. 1986. *Parallel*

distributed processing: explorations in the microstructure of cognition. Cambridge, Mass: MIT Press.

Snoad, N. and Bossomaier, T. 1995. MONSTER - the ghost in the connection machine: modularity of neural systems in theoretical evolutionary research. http://www.chg.ru/SC95PROC/531_NSNO/SC95.HTM.

Stevens, C. 1994. Complexity of brain circuits. In Cowan, G. A., Pines, D., and Meltzer, D. eds. *Complexity: metaphors, models and reality*, 245-261. Reading, Mass.: Addison-Wesley.

Steward, J. 1997. The evolution of genetic cognition. *Journal of Social and Evolutionary Systems* 20:53-73.

Tautz, D. 1992. Redundancies, development and the flow of information. *BioEssays* 14:263-266.

Wagner, G. P. 1995. Adaptation and the modular design of organisms. In Moran, F., Moreno, A., Merelo, J.J., and Chacon, P. eds. *Lecture notes in artificial intelligence: advances in artificial life*, 317-328. Berlino-Heidelberg: Springer-Verlag.

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

Wagner, G. P., 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.