

On the Modelling of Dynamical Hierarchies: Introduction to the Workshop WDH 2002.

Tom Lenaerts¹, Dominique Groß² and Richard Watson³

¹Computational Modelling Lab, Vrije Universiteit Brussel, Belgium

²Senter for Vitskapsteori, University of Bergen, Norway

³Organismic and Evolutionary Biology Department, Harvard University, USA

Abstract

This paper provides an introduction to the first workshop on the modelling of dynamical hierarchies. We focus on a number of relevant issues including the difference between structural and dynamical hierarchies, different biological and artificial metaphors related to this subject and a small set of attempts to design software artifacts which build on these notions. We conclude with questions for discussion.

Introduction

When observing natural systems, hierarchical composition is ubiquitous in both the physical and the biological. On the one hand, the physical hierarchy from subatomic particles, to atoms, to molecules, through to planets, solar systems, and galaxies has a composition that is familiar and that serves in making models of the rest of the world. On the other hand, the levels of organisation in biological systems from organic molecules, to cells, to organs, to organisms, through to populations and ecosystems are also commonplace tools in understanding the biota.

Apart from this observational fact that hierarchies are present in many different natural systems and that they assist us in understanding these systems, other motivations can be listed for studying this topic: 1) Through the study of the structuring process we can maybe understand why it is so pervasive and 2) Through the design and development of computational models which mimic the recursive emergence of these hierarchies we can better understand how they are formed (i.e. the problem of synthesis) and can be understood in terms of their qualitative dynamics. This is also still an important issue in Biology (31).

This paper and the associated workshop specifically focuses on the second motivation. The issue of modelling the synthesis of global phenomena from local interactions between elementary units through agent-based models is central to Alife research (10; 26). Yet, it is not the intention of the workshop to cover the entire spectrum of

Alife research. It will mainly focus on four questions: 1) How can hierarchies, as can be observed in biological and physical systems, emerge recursively through the dynamical interactions of simple and aggregate components, 2) Which metaphors from Biology or Physics are suitable as guidelines to build these models, 3) How can we formally define the entire process in order to achieve a general definition and 4) What is the relevance of these models in the context of other research domains or engineering applications?

This paper provides an introduction to the workshop and highlights those issues which are considered relevant. It is not our intent to be comprehensive in reviewing the literature on the subject of dynamical hierarchies, nor do we pretend to resolve the deep conceptual issues involved in this subject. Instead, at this stage, we intend to use this workshop paper only to promote discussion of the topic, to outline some connections across disciplines and to present our thoughts.

In the next two sections, two examples of hierarchically structured systems will be given. These examples will try to provide a clear distinction between a *structural* hierarchy and a *dynamical* hierarchy. In the next section, the metaphors which are often used to elaborate on the dynamical hierarchies are discussed. Furthermore specific attention will be paid to three topics from biology which we consider important: multi-level selection, symbiogenesis and evolutionary transitions. Afterwards the relevance of dynamical hierarchies in the context of problem solving is discussed. Finally, a small discussion is provided on the recent attempts that were made to define higher-level structures and open questions which should be addressed are listed.

The reader should be aware that some of these notions and concepts are sometimes preliminary and that there exist different schools of thought on the different topics.

Separating Structural and Dynamical Hierarchies

Most Alife researchers are familiar with designing software artifacts in some favourite programming language. It is considered bad practice, in terms of re-usability and maintainability, to write these programs as a giant web of interconnected procedures and functions. Hence, a number of design heuristics have been introduced to improve application code.

One of these directives stresses that one has to write programs using *layers of abstraction*. This means that programmers should write their code in modules which are connected in a hierarchical fashion: from the programming language (at the lowest level) until the final application (at the highest level). The functions at each new layer are recursively composed of functions and/or primitives defined at the lower layers. Hence, there is a vertical coupling between different functions. Furthermore, each function defined at a particular level is considered modular in the sense that the body of these functions can be changed or improved independently of the other functions at the same level. Note, that the resulting hierarchy is not necessarily pyramid-shaped. A limited number of functions at the lower level can be used to define an exponential number of functions at a higher level. An example is the definition of a new language on-top of another programming language (like Lisp).

The distinction between structural hierarchies, implicit in the programming example, and dynamical hierarchies, is due to the difference in mechanisms which produce the hierarchy. In the programming example, the hierarchy was formed by imposing a part-whole relationship on the programs, subprograms, functions and primitives. Hence, it was engineered through the separation of different parts into modules and abstraction layers. In the next example the hierarchy is the result of the interaction of different lower-level components. These interactions produce new structures which obtain their own functionality and mechanisms to form new higher structures.

Proteins are a very important and powerful class of macromolecules. They are generally composed of a linear sequence of a few hundred amino acids that are of less than two dozen types. Either during or after synthesis, the protein folds into a secondary and tertiary structure which defines both its spatial form and its biochemical properties. In organisms, the functions of proteins usually involve interactions with ions, small molecules or macromolecules and include most importantly catalytic functions of high specificity. Also very important are different binding and splicing and cutting operations on DNA and RNA and other proteins. The interplay of the multitude of types of proteins (generally thousands) synthesized in cells is therefore absolutely necessary, albeit

not sufficient, for the complexity of living systems.

Note that although proteins acquire their properties through the interactions of their components, the amino acids (which themselves have specific chemical properties), these properties are largely due to their overall structure. Like the modular functions in the programming example, proteins are highly complex structures that consist of relatively simple building blocks. Yet the properties and functionality of the module are a result of the overall structure of the module and not so much the properties of each specific component.

Proteins themselves are building blocks of higher order entities (organisms) which in turn ensure the production and maintenance of the proteins. This catalytic or causal closure is sometimes seen as a defining characteristic of life itself (13). The relevant keyword in this context is *autopoiesis* (38).

Contrary to the structural hierarchy example, proteins are members of a dynamical hierarchy: If mixed together they will typically interact in a number of interesting ways, change their internal configurations in reaction to external stimuli, decay, unfold, re-fold, form new proteins or catalyse reactions or display other dynamic behaviours. Proteins, can therefore be regarded as *autonomously dynamic* agents. Finally, note that dynamical hierarchies are in no way restricted to living systems.

Concepts and Metaphors

When reasoning about dynamical hierarchies, of which the earlier described proteins are a member, different intuitive notions have been examined. One of these notions is the concept of nearly-decomposable systems described by H. Simon in (33). In this section we will discuss this concept and the related concept of modularity. In the context of dynamical hierarchies, it seems that the ideas of nearly-decomposable systems, loose coupling and modularity describe in a clear way our intuitive notions of this topic. Afterwards, we will discuss the relationship of these concepts with three biological metaphors which we consider relevant: natural selection (especially multi-level selection), symbiogenesis and evolutionary transitions.

Nearly-Decomposable Systems

Based on the discussion of a simple parable of two watchmakers, Simon deduces that for complex systems to evolve through natural selection, some form of modularity in the components, which constitute the system, seems to be important to increase the probability that the system can be constructed (33). Simon described two watchmakers who, starting from the simplest components, try to build a complete watch. The first watchmaker does this by incrementally adding one piece after

the other. The second watchmaker first starts by combining primitive components into small modules which are then combined again forming more complex modules and so forth. However, both watchmakers are interrupted in their work at regular intervals, which results in the disintegration of the module they were working on. Now, if an interruption always occurs during the construction of a watch, the first watchmaker will never be able to build any watches since they will always disintegrate as soon as he is interrupted.

Simon uses this parable to illustrate that the time required for the evolution of a complex form depends critically on the numbers and the distribution of potential intermediate stable forms (modules). Concretely, this means that, under the assumption that there exists such a hierarchy of modules and the modules at each level in this hierarchy have the same *span* (i.e. the number of components which constitute the higher-level module) then the time required to build, for instance, a procedure from primitives is the same as the time required to build a program which consist of the same amount of intermediate procedures. If there does not exist such a hierarchy of stable modules, the chance of a very complex program being generated from scratch will be highly improbable.

It is important to see that, in the simple metaphor of the two watchmakers, the selection process is not provided by some vitalistic external force (the watchmakers), but solely through the stability of the complex forms which arose through the *interaction* of simple ones, i.e. survival of the stable. Hence, the relation with the dynamical hierarchies concept: higher-order structures are a product of the interactions and properties of the lower-level components.

Furthermore, the interactions which occur in these modular hierarchic structures can occur *within* the different modular entities and *between* the different modules. The timings of these interactions at the different levels are considered to be of a different order of magnitude. This property of nearly-decomposable systems is directly applicable to biological systems: The intra-cellular dynamics of one cell generally act at a faster time scale than the inter-cellular interactions, and they are largely independent of the detailed intra-cellular interactions within another cell.

A system can be partitioned in different ways. When there is a vertical separation between the slower higher level, the faster lower level dynamics, and the dynamic at a particular focal level which is under investigation, the systems is referred to as a *decomposable* system. In such a system where the higher-level and/or lower-level dynamical influences are weak but not negligible, it is called a *nearly-decomposable* system. This kind of separation is often performed when determining a system of differential equations for some natural phenomenon.

The separation of a decomposable system can also be extended to the separation between horizontal modules at the same hierarchical level. An example of the latter is the geographical separation between colonies of a similar breeding species. This kind of system is nearly-decomposable when individuals of this species are still able to migrate between the different colonies. Nearly-decomposable systems have two important properties: 1) the short-run behaviour of each of the component sub-systems is approximately independent of the short-run behaviour of the other components, and 2) in the long run the behaviour of any one of the components depends in only an aggregate way on the behaviour of the other components (33). A further discussion of this separation will be addressed in the next section (see also Watson 2002 in this volume).

Note that in some biological instances of these nearly-decomposable systems selective forces can still act on the different entities and aggregated entities in the environment, since interaction between entities can be competitive and hence lead to different reproduction rates.

Modularity

Simon's characterisation of nearly decomposable systems is as much a characterisation of modularity as it is a characterisation of hierarchy. Notions of hierarchy and modularity seem to be closely related in the sense that when a whole system is understandable as the interaction of a number of sub-systems then we necessarily have at least a two layer hierarchical system. Whereas, when a system is not modular and we are unable to discern structure at a sub-system scale, there seems to be little value in describing a system as having more than one level of organisation.

Modularity refers to a property of a system and can be observed when there exists a *loose horizontal coupling* between the entities at the same level (34). Loose horizontal coupling between assemblies at the same level allows each assembly to operate independently of the details of the other assemblies at the same level. Hence, only the inputs and outputs of each assembly are important for the dynamics at that level.

In a nearly-decomposable system, *loose vertical coupling* is a second important property which affects the dynamics in the structure. This is exactly the property of a system which produces the hierarchical nature of the system. This second form of loose coupling allows levels to be identified in a natural system. As a result the dynamics at the particular levels can be described almost independently of the dynamics of the lower-level entities which constitute this level and the higher-level entities it produces. Note though, that in order to allow interaction between levels, some form of filtering and buffering between these levels is necessary (17).

In a biological context modularity can be an important property of the genotype-phenotype map. As defined by Wagner and Altenberg (39), a genotype-phenotype mapping can be considered modular, if there exists an independent genetic representation of different traits. In a modular genotype-phenotype map there are few pleiotropic effects among genes serving different functions. The pleiotropic effects are clustered mainly among genes that are part of a single functional complex or trait.

Simon provides an abstract example which presents a clear notion of the benefits of a modular system (33). Imagine a lock which consists of 10 dials each with 100 settings. How long will it take to open the lock if every combination of the 10 dials has to be tried and how long will it take to open the lock when we can set each dial separately, i.e. independent of the settings of the other dials? The latter question refers to a system which consists of 10 independent modules as opposed to 1 big module in the first case. In answer to this question it should be clear that the latter scenario corresponds to the modular structure where only $10 * 100$ combinations need to be tried instead of 10^{100} .

From an evolutionary perspective, modularity is important since it allows mutation and selection to affect one feature of an organism without interfering with the others (16). As a result, the adaptation of one functional property does not influence the other properties/phenotypic traits of the organism.

Although our intuitive picture of dynamical hierarchies has become a little clearer after describing these concepts, there is still no formal definition which clearly describes the sufficient and necessary conditions to identify a truly dynamical hierarchy. Hence, we require more information in order to finally derive these conditions. Therefore, in the next sections, we will discuss a number of issues in Biology which we think, can improve our understanding.

Modularity and Selection

The concept of a module is critical in discussing issues related to evolution through natural selection.

Evolution by means of natural selection depends centrally on heritable variation and differential reproduction. These properties both depend on an ability to conceptually separate a biological entity from its environment. Intuitively, the very existence of a coherent entity depends on its internal stability and (semi) independence from its environment as is often taken for granted when thinking about evolution. Hence, there exists a strong relation between the notion of selection and modularity.

Speaking in general, if these modules, sometimes referred to as vehicles or interactors, are capable of reproduction in some way with the possibility of introduc-

ing variation and this reproduction process depends in some way on the phenotypic aspects of these entities, then these entities can be considered to be *units of selection*. The stable *units* can interact and compete with each other in their particular environment. As a result of this competition differential survival based on phenotypic traits occurs, i.e. natural selection.

Now, how does natural selection relate to the dynamic formation of hierarchical molecular structures? We already know, as was discussed in the subsection on modularity, that, in terms of selection, modular genomes allow the adaptation of features without any interference within other features. But these modular structures, in addition to enhancing evolvability when present, are themselves a product of evolution. What is the impact of selection on this formation?

E. Szathmary describes this issue in the context of early (primitive) biological systems (37). In these systems, the first nucleic acids, primitive carriers of genetic information, could not have been very long since, in the absence of specific replicases, copying would have been inaccurate. As a result, large molecules accumulated errors. Thus, primordial genomes must have consisted of smaller pieces of nucleic acids. This raises an important question: How could cooperating groups of small replicators have arisen, and how could they have protected themselves against invading molecular parasites.

In answer to this question, Szathmary lists two phenomena that underlie all increases in complexity and cooperation in evolution (37). The first phenomenon, called *synergism*, specifies that the combination of different functions can result in strong synergistic fitness effects which can allow these complex entities to survive. The second phenomenon is *limited dispersal* and is the result of localisation and/or encapsulation. Hence, this phenomenon results in the dependence between different individuals and their offspring in future generations. This dependence results from the facts that they can be considered to be in the same boat, i.e. one can not survive or reproduce without ensuring the survival or reproduction of the others.

Note that this question is also related to a different problem called the Eigen's paradox (37). It is assumed that the copying process of larger genomes results in more errors. To improve the fidelity of the copying process, particular enzymes are required. Yet, the genomes which code for such enzymes are large. Due to this property they again require enzymes to ensure the fidelity of their a copies. Hence, a paradox occurs. Eigen attempted to resolve the paradox and therefore introduced the concept of hyper-cycles, i.e. a self-producing system consisting of several smaller independently self-replicating systems. Although these hyper-cycles are able to form they are still not stable since they can eas-

ily be destroyed through other parasitic individuals in the environment. So the question, in this case is, how can networks of interacting (collaborating) components evolve and survive in a natural (or artificial environment).

Multi-level Selection in Nearly-Decomposable Systems

The issue of how cooperative groups of small replicators have arisen is directly related to a fundamental question posed by E.O Wilson (also called the central theoretical problem of sociobiology, (42)): How can altruism which by definition reduces personal fitness, possibly evolve by natural selection? Different attempts have been made to answer this question; e.g. Evolutionary Game Theory and Kin selection theory (7; 8; 23; 35).

The general idea is that in order for complex networks of simple particles to evolve through natural selection, the network has to possess beneficial properties which allow it to survive in its competitive context.

From the research produced by sociobiologists it became clear that the altruistic trait can evolve under specific conditions posed on the structure of the population. Important models which incorporate this notions are: the trait-group selection model by Wilson and its extension by Michod (24; 35), the haystack models proposed by Williams and Williams and also by Maynard Smith (35; 2) and for the hyper-cycle problem the Stochastic Corrector Model proposed by Szathmary (24; 37). The importance of these models lies in the fact that they provide clues on how to construct selectionist models which allow for the emergence of dynamical hierarchies.

The dynamics of selection in these models can be analysed using for instance replicator dynamics from evolutionary game theory (41) or the Price covariance equation (28; 29). Replicator dynamics are dynamical systems which highlight the role of selection in an evolutionary model. An analysis of a simple example is provided by Lenaerts in (18). The Price covariance equation is often cited in the context of multi-level selection because it is a scalable equation since it can be formulated in terms of any quantitative property. Price expresses the rate of change of a particular quantitative property in a population as the combination of the intensity of selection on that quantity and the heritable variation which is available for selection. This expression can be written for any quantitative property, i.e. for individuals within a population and for modules (which consist of individuals) in another population.

The selection dynamics, expressed by both equations, within each module are usually considered to be stronger than selection between modules since there is more variation within than between these modules. Nevertheless

under certain conditions the balance between higher and lower-level forces can change from the first to the second and vice versa (18). These conditions, as discussed before, are related to localisation or encapsulation. More information can also be found in (11).

Symbiotic Relationships

A direct example of the phenomena that Szathmary holds responsible for the increase in complexity refers to a more radical process of interspecific genetic integration. Endosymbiosis, for example, is the symbiosis (living together) of two organisms where one lives inside the other. In some cases of endosymbiosis, it is very difficult to establish the independent status of the parties involved.

Indeed, the endosymbiosis of prokaryotes (bacteria) with one another has to a large extent become irreversible, i.e. the parties involved have subsequently become dependent on one another and cannot live independently. Such endosymbiosis accounts for the origin of various organelles, e.g. mitochondria and chloroplasts, in eukaryote cells. And there is some evidence to suggest that other organelles, such as flagella, also resulted from the physical inclusion of symbiotic bacteria (20).

In some cases, the symbiotic entities remain as physiologically distinct *organelles* within the larger host cell, e.g. plastids such as mitochondria and chloroplasts. But there are also mechanisms, such as transposition, that can transfer genetic material from plastids to the nucleus or vice versa, implying that symbiotic origins may account for nuclear DNA as well as cytoplasmic DNA, as Doolittle's data concord (5).

Such evidence suggests a more general process *symbiogenesis*, the origin of new species from the integration of symbionts. In this view, endosymbiosis is not simply responsible for the origin of eukaryote organelles, but for the origin of the eukaryotes themselves (including all plants and animals).

Margulis describes the process as follows:

[...] different bacteria form consortia that, under ecological pressures, associate and undergo metabolic and genetic change such that their tightly integrated communities result in individuality at a more complex level of organization [...] (22).

Moreover, evidence suggests that secondary endosymbiosis has also taken place; that is, endosymbiosis involving the inclusion of an endosymbiont. This process of creating new species from symbiotic relationships is known as symbiogenesis, and is the basis of Serial Endosymbiosis Theory (19; 21). This theory proposes that endosymbiosis has been a recurring source of innovation in the evolution of complex cells.

The mechanisms and effects of symbiogenesis are under-researched and currently poorly understood. Moreover, the notions of serial endosymbiosis suggest an idea of repeated inclusion and hierarchical encapsulation, the consequences of which are even more poorly understood. Hierarchical encapsulation is one of the features common in *the major transitions in evolution* (discussed below).

The existence of endosymbiosis, symbiogenesis, and the major evolutionary transitions causes us to accept that the generally accepted levels of encapsulation from genes to cells to organisms should not be taken for granted. Furthermore, the creation of new levels of organization has been absolutely fundamental in evolutionary history.

The Major Evolutionary Transitions

The major evolutionary transitions (3; 24; 25) involve the creation of new higher-level complexes of simpler entities. Hence these theories, capture aspects of the previous two sections which examined models for the increase of complexity in biological systems.

Summarised by Michod for example, they include the transitions from individual genes to networks of genes, from gene networks to bacteria-like cells, from bacteria-like cells to eukaryotic cells with organelles, from cells to multicellular organisms, and from solitary organisms to societies (25). There are many good reasons to be interested in the evolutionary transitions: they challenge the Modern Synthesis preoccupation with the individual as the unit of selection, they involve the adoption of new modes of transmitting information, and they address fundamental questions about individuality, cooperation, fitness, and not least, the origins of life.

In several of the transitions, entities, that were capable of independent replication before the transition, can replicate only as part of a larger whole after it (24; 25). Although Maynard Smith and Szathmary identify several transitions which do not fit what they describe as symbiosis followed by compartmentation and synchronised replication, several of the transitions including the origin of eukaryotes from prokaryotes (21; 20), and the origin of chromosomes from independent genes (24), do involve the quite literal union of previously free-living entities into a new whole.

This form of change has some obvious contrasts with how we normally understand the mechanisms of neo-Darwinist evolution. The ordinary (non-transitional) view of evolutionary change involves the accumulation of relatively small random genetic variations within an entity, whereas innovation by composition involves the union of two different entities, each contributing relatively large amounts of pre-adapted genetic material to the new composite entity (40).

Modularity and Hierarchies in Machine Learning

Although there exists a reasonable body of biological theories on how complexity emerges in an evolutionary context (see previous section), it is still difficult to design and develop an artificial system which possesses the same properties and mimics the dynamical construction of modular hierarchies. What are the properties of the essential particles and how do they interact? What defines a unit at a higher abstract level and what are the ‘emerged’ properties at that level? What new forms of interaction emerge at the next level? If we move upwards in the hierarchy, does variation in terms of the number of different ‘units’ at each level decrease or increase?

The answers to these questions are in the first place important to get a better understanding of the workings of the different theories. Furthermore, the answers are also important outside the context of strictly modelling biological and physical phenomena. For instance in the domain of Evolutionary Computation (EC) and more generally Machine Learning (ML) the resulting information can also prove important.

There are a number of artificial evolution methods in EC that seek to explicitly represent and exploit modularity. Notably, modularity is addressed *implicitly* by the use of variable-length, moving-locus, non-linear, and generative encodings - for example, linkage learning genetic algorithm (9). Moreover modularity is addressed *explicitly* in mechanisms that encapsulate subsets of features for subsequent re-use during the search process - for example, adaptive representation (32).

Methods explicitly addressing problem decomposition (the division of labour) by using selection on *parts* as well as the *whole* solutions include Learning Classifier Systems (15), Cooperative Coevolution (27). These methods indicate a broad interest in artificial models of evolution that explicitly utilise the aggregation of smaller units into larger wholes. However, there have been very few examples that illustrate the multi-level assembly of parts into wholes over successive hierarchical levels, e.g. Compositional Evolution (40).

Furthermore, in ML, different techniques such as *hierarchical learning* and *layered learning* were constructed to improve the generalisation of a single learning task or combine the effects of different sub-tasks (36). In the first technique, the idea is that the combined result of a number of generalisers which work on the same problem gives a better generalisation of the learned task as compared to the results of the generalisers in isolation. Layered learning is different in the sense that each generaliser has learned a different sub-task.

Especially in the niche of reinforcement learning, researchers are very interested in hierarchical learning approaches. In order to achieve that, the algorithms incor-

porate *gated* behaviours, i.e. there is a set of behaviours which act when the environment is in a certain state and a gating function that decides which action should be selected and actually executed. In some cases only the behaviours are learned, in other cases the gating function and again in other cases both of them (12).

Note, that in the hierarchical learning and hierarchical reinforcement learning approaches, the task composition is all done manually. Up until now, there only has been one attempt, we know of, to learn the task decomposition. This was done using the Nested Q-learning algorithm (4). This algorithm generates its own hierarchical control structure and then learns the low-level skills together with the function which will select among them.

Questions for Discussion

We have examined a set of concepts and biological metaphors which can aid us in the design of an agent-based model for the dynamic construction of a hierarchy. We have seen that the research is relevant from both the viewpoint of understanding natural systems and the application of the resulting models in other domains. Yet in everything we have described so far there is a lack of formal specification. How are higher-level modules defined? How do we distinguish between inherited properties on the one hand and emergent properties on the other hand? How do we evaluate any agent-based model which claims to produce higher-level structures through the dynamical interactions of lower-level entities. Hence, a definition is required to outline under which conditions these structures can be called real higher-order structures. These questions have to be answered since the resulting answers would provide useful information for a decent definition of the necessary and sufficient conditions of a dynamical hierarchy. As a result we would no longer need to rely on an intuitive notion to validate the claim that a particular agent-based model mimics a dynamical hierarchy.

The importance of this issue has already been put forward by several researchers. In (1), Baas introduces a general framework for the study of emergence, hierarchies and hyper-structures. This framework tries to unify the different definitions of, and approaches to the study of, emergence and hierarchies using the notion of hyper-structures. This new notion provides us with a tool to describe a complex system in a reasonable way.

Based on this notion Rasmussen and coworkers' (30) attempt to demonstrate the spontaneous emergence of micelles from the bottom-up. In their model they are able to show how primitive agents (the "monomers") interact with one another to form polymers (claimed to be second order structures). Those polymers then in turn assemble into micelles (claimed to be third order structures). There has been some recent controversy concern-

ing the exact interpretation of the hierarchical order of those micelles (6).

Although Rasmussen's micelles certainly satisfy the intuition about what a third order structure should be, it remains only intuitive. A more rigorous criterion, that specifies under which condition a composite object forms a new hierarchical layer, is therefore indispensable. Rasmussen and coworkers, already try to provide this. Furthermore, this criterion should be validated on the different relevant biological theories to make sure it captures the most important facets.

These questions and others, already listed in the rest of the paper, are the topic of the workshop. It is, in our opinion, only through an open discussion, that certain problems can be resolved and our intuitive notions about all these concepts can crystallize into concrete forms.

Acknowledgments

We want to acknowledge the organisers of the eighth ALife conference who allowed us to organise this workshop. Many thanks also to the members of the program committee: Nils A. Baas, Mark Bedau, Larry Bull, Claus Emmeche, Inman Harvey, Paulien Hogeweg, Barry McMullin, Bernard Manderick, Victor Norris, Norman Packard, Steen Rasmussen, Stanley N. Salthe, and Eörs Szathmáry. We are also in debt to Roger Strand for his advice on this paper and Ann Nowé for doing some last minute reviewing.

References

- N. Baas. Emergence, Hierarchies, and Hyperstructures. In C. Langton, editor, *Artificial Life III*, volume XVII of *Studies in the Sciences of Complexity*, pages 515–537. Santa Fe Institute, 1992.
- T.D. Bergstrom. Evolution of social behaviour: Individual and group selection models. Paper provided by California Santa Barbara - Department of Economics in its series University of California Santa Barbara - Department of Economics with number 14-01., 2001.
- L.W. Buss. *The Evolution of Individuality*. Princeton University Press, New Jersey, 1987.
- B. Digney. Emergent hierarchical control structures: Learning reactive/hierarchical relationships in reinforcement environments. In *Proceedings of the 4th International Conference of Simulation of Adaptive Behavior*, pages 363–372. MIT Press, 1996.
- W.F. Doolittle. Uprooting the tree of life. *Scientific American*, 282:90–95, 2000.
- D. Groß and B. McMullin. Is it the right ansatz. *Artificial Life*, 7(4):355–366, 2001.
- W.D. Hamilton. The genetical evolution of social behaviour i. *Journal of Theoretical Biology*, 7:1–16, 1964.

- W.D. Hamilton. The genetical evolution of social behaviour ii. *Journal of Theoretical Biology*, 7:17–52, 1964.
- G.R. Harik and D.E. Goldberg. Learning linkage. In R.K. Belew and M.D. Vose, editors, *Foundations of Genetic Algorithms 4*, pages 247–262. Morgan Kaufmann, San Mateo CA., 1996.
- P. Hogeweg. As large as life and twice as natural: Bioinformatics and the artificial life paradigm. *Complexity - From Biology to Computation*, pages 2–11, 1993.
- P. Hogeweg. Multilevel evolution: replicators and the evolution of diversity. *Physica D*, 75:275–291, 1994.
- L.P. Kaelbling, M.L. Littman, and A.W. Moore. Reinforcement learning: A survey. *Journal of Artificial Intelligence Research*, 4:237 – 285, May 1996.
- S. Kauffman. *Origins of Order : Self-organization and Selection in Evolution*. Oxford University Press, 1993.
- L. Keller, editor. *Levels of Selection in Evolution*. Monographs in Behaviour and Ecology, Princeton University Press, 1999.
- P.L. Lanzi, W. Stolzmann, and S.W. Wilson, editors. *Learning Classifier Systems; From Foundations to Applications*. Lecture Notes in Artificial Intelligence; Vol. 1813; Subseries of Lecture Notes in Computer Science. Springer, 2000.
- E.G. Leigh, Jr. *Levels of Selection, Potential Conflicts, and Their Resolution: The Role of the "Common Good"*, pages 15–30. In Keller (14), 1999.
- J. Lemke. Opening up closure: semiotics across scales. *Annals New York Academy Sciences*, 901:100–111, 2000.
- T. Lenaerts, A. Defaweux, P. van Remortel, and B. Manderick. Multi-level selection in a simple evolutionary model. In Hendrick Blockeel and Marc Denecker, editors, *Proceedings of the Fourteenth Belgium-Netherlands Conference on Artificial Intelligence*, pages 203–210, 2002.
- L. Margulis. *Origin of Eukaryotic cell*. Yale University Press, New Haven, 1970.
- L. Margulis. Origins of species: Acquired genomes and individuality. *Biosystems*, 31(2-3):121–125, 1993.
- L. Margulis. *Symbiosis in Cell Evolution; Microbial Communities in the Archean and Proterozoic Eons*. WH Freeman and Co. New York, 1993.
- L. Margulis. Gaia is a tough bitch. In J. Brockman, editor, *The Third Culture: Beyond the Scientific Revolution*. Simon and Schuster, 1995.
- J. Maynard Smith. *Evolution and the Theory of Games*. Cambridge University Press, Cambridge, UK, 1982.
- J. Maynard Smith and E. Szathmáry. *The major transitions in evolution*. Oxford University Press, 1995.
- R.E. Michod. *Darwinian Dynamics; evolutionary transitions in fitness and individuality*. Princeton University Press, New Jersey, 1999.
- N.H. Packard and M.A. Bedau. Artificial life. In *Encyclopedia of Cognitive Science*. Macmillan Reference Ltd, 2000.
- M.A. Potter. *The Design and Analysis of a Computational Model of Cooperative Coevolution*. PhD thesis, George Mason University, Fairfax, Virginia, 1997.
- G.R. Price. Selection and covariance. *Nature*, 227:520–521, 1970.
- G.R. Price. Extension of covariance selection mathematics. *Annals of Human Genetics, London*, 35:485–490, 1972.
- S. Rasmussen, N. Baas, B. Mayer, M. Nilsson, and M. Oleson. Ansatz for dynamical hierarchies. *Artificial Life*, 7(4):329–353, 2001.
- C. Ripoll, J. Guespin-Michel, V. Norris, and M. Thellier. Defining integrative biology. *Complexity*, 4:19–20, 1998.
- J. Rosca. *Hierarchical Learning with Procedural Abstraction Mechanisms*. PhD thesis, University of Rochester, Rochester, New York, 1997.
- H.A. Simon. *The Sciences of the Artificial*. The MIT Press, Cambridge, Massachusetts, London, England., 1969.
- H.A. Simon. The organization of complex systems. In *Hierarchy Theory: The Challenge of Complex Systems*, pages 1–27. 1973.
- E. Sober and D.S. Wilson. *Unto Others, the Evolution and Psychology of Unselfish Behaviour*. Cambridge, MA:Harvard University Press, 1998.
- Peter Stone and Manuela M. Veloso. Layered learning. In *Machine Learning: ECML 2000, 11th European Conference on Machine Learning, Barcelona, Catalonia, Spain, May 31 - June 2, 2000, Proceedings*, volume 1810, pages 369–381. Springer, Berlin, 2000.
- E. Szathmáry. *The First Replicators.*, pages 31–52. In Keller (14), 1999.
- F.J. Varela. *Principles of Biological Autonomy*. The North Holland series in general systems research. North Holland, New York, 1979.
- G.P. Wagner and L. Altenberg. Complex adaptations and the evolution of evolvability. *Evolution*, 50:967–976, 1996.
- R.A. Watson. *Compositional Evolution: Interdisciplinary Investigations in Evolvability, Modularity and Symbiosis*. PhD thesis, Computer Science Department, Brandeis University, MA, USA, 2002.
- J. Weibull. *Evolutionary Game Theory*. MIT Press, Cambridge, MA, 1996.
- E.O. Wilson. *Sociobiology: The New Synthesis*. The Belknap Press of Harvard University Press, Cambridge, Massachusetts and London, England, fourth edition, 1976.