From Directed to Open-Ended Evolution in a Complex Simulation Model

Maciej Komosiński[1] and Ádám Rotaru-Varga[2]

[1]Institute of Computing Science,
Poznan University of Technology
Piotrowo 3A, 60-965 Poznan, Poland

[2]School of Computing Science
Simon Fraser University
Vancouver, Canada

Maciej.Komosinski@cs.put.poznan.pl

http://www.frams.poznan.pl/

Abstract

The problem of achieving open-ended evolution in complex systems is studied in this paper. We propose various techniques that support it, and present a gradual approach. These techniques are used and tested in the *Framsticks* system, which is a realistic, three-dimensional artificial life simulator with rich capabilities. Specifically, as there are no strong constraints imposed on the structure (body) and control (neural network brain), this system is suitable for testing open-endedness. In Framsticks, energetic requirements form the basis for competition, while interactions occurring on various levels act as a source of complexity and variation. The way towards open-ended evolution is discussed, developmental genotype encodings are proposed, and the results of so-far experiments are presented.

Introduction

Those of us who have ever run an artificial evolution experiment probably know the excitement of observing evolution unfolding under our guidance, and arriving at the solutions we imagined. Even bigger is our excitement when evolution surpasses our initial expectations, producing results never imagined. More exciting as it is, spontaneous, open-ended evolution rarely occurs. There are several theoretical works investigating reasons for this phenomenon and providing insights how to approach it. There are also several practical works focusing on *open-ended* evolution using relatively simple models (Adami and Brown 1994, Balkenius 1995, Channon and Damper 1998). In this work, we analyze such evolution in a complex, realistic system. We present theoretical arguments, discuss why open-endedness is important to study, why it is hard to accomplish, and what are some of its prerequisites. The gradual approach to achieving open-endedness is also described.

Our system, *Framsticks*, was implemented with several hierarchical objectives in mind: a fast simulation of a complex, realistic 3D world; autonomous creatures with free-form body and brain; ability to design and test creatures interactively; possibility of evolution of creatures to solve predefined problems (locomotion, swimming, food-finding, etc.); and ultimately, *spontaneous evolution of complex creatures* and behaviors. The development of Framsticks (Komosiński and Ulatowski 1999, Komosiński 2000) has satisfied these objectives, except for spontaneous evolution, which is our current focus. The property of open-endedness in such an environment may bring meaningful results, with potential references to natural evolution and the real world.

The creatures in Framsticks are composed of connected 'sticks', equipped with muscles, sensors and a neural network. They are capable of interacting with the environment (walking, swimming, ingesting), and each other (locating, pushing, hurting, killing, eating, etc.). The virtual world can contain a combination of flat land, hills, water, and various objects. Despite the simplified simulation (made possible by the sticks), this setup allows for complex behaviors.

Evolution operates on a population of genotypes (describing brains and bodies), using mutation and crossover. Several predefined criteria can be used to direct evolution (fitness). A host of interesting behaviors (Komosiński and Ulatowski 1997) is a result of this technique (some of them are reported here). Using a general criterion for selection, such as life span, favors reproduction of organisms which can survive efficiently, but does not restrict them to specific behaviors. This is one of our techniques for approaching open-ended evolution. Other proposed techniques include coevolution and natural selection.

Potentially, as there is no limitation imposed on the complexity of neural networks, creatures can display sophisticated behaviors. They can sense (and possibly remember) their environment, find out their orientation, locate food sources and other creatures. With the proper use of sensors and complex brains, creatures could discriminate between various morphologies of other habitants, sense their movement, exhibit preferences, and various group and social behaviors. The simulation of the virtual environment allows for *embodiment of intelligence*, which is argued to be an important factor when studying intelligence (Hofstadter 1979, Bajcsy 1998).

The following section describes our ideas on approaching open-endedness in a complex system. Section III presents the Framsticks system in more detail, and section IV characterizes the proposed genetic representations. Finally, we describe some of our evolutionary experiments, and outline future work.

Towards Open-Ended Evolution

As evolution in artificial systems had materialized in numerous works in the past decade, the distinction between *directed vs. open-ended* evolution was identified by several authors. This theme, referred to as "spontaneous evolution", "incremental adaptation" (Cliff, Harvey and Husbands 1993), or "perpetuating evolution" (Channon and Damper 1998), has been described from several perspectives, often independently. As the concept of open-ended evolution does not have well-defined boundaries, we are prompted to present our interpretation.

Intuitively, open-ended evolution happens if a system continues evolving. A more rigorous formulation is as follows: a system displays open-ended evolution when some internal complexity of the system can (and does) continually increase, with minimal outside involvement by the designer of the system. A more practical question is how to design and build such a system. Based on various suggestions and results, we compiled a list of necessary conditions (which does not have to be sufficient):

- the simulation model is not limited in complexity,
- interactions occurring on different levels act as a source of complexity,
- evolution works on highly evolvable genetic representations.

Since Framsticks satisfies mostly the first two requirements, our current focus is on more evolvable genotype representations. The internal complexity of creatures and inter-creature interactions are potentially unlimited. There are several types of other interactions: between parts of a creature, between creatures and environment, between groups of creatures (species), etc.

As we achieved interesting results with Framsticks with less and less human intervention, we learned that there are *degrees* to open-endedness. At one extreme, the designer's involvement can be total, when every detail is designed directly (this case is more interesting from an engineering standpoint than from artificial life). Evolution is given some role when it is used to improve hand-built creatures. When evolution is used to create creatures that satisfy some predefined criteria, human involvement is still present, but on a different level. Finally, in an open-ended scenario, the designer's involvement is minimal, and potential goals of evolution are not known *a priori*.

Based on this succession of scenarios, we feel that there is a gradual approach to open-endedness: it is not an all-or-nothing property. True open-endedness is difficult to achieve, but we believe Framsticks advances in the right direction.

Framsticks Ecosystem

System Architecture

Framsticks creatures live in a virtual world. However, a distinction is introduced between the gene pool and the living phenotypes. The capacity of the world and the capacity of the gene pool can be adjusted, thus constituting a simulation ratio. When the interactions between creatures can be ignored, only one phenotype may be simulated at a time. On the other hand, when the interaction is important and coevolution is to take place, the number of concurrently living organisms may be high, possibly equal to the capacity of the gene pool. The more individuals are simulated, the stronger is the competition between species (groups of individuals which exhibit similar methods of achieving fitness, for example, by acquiring energy).

Such separation of the pool of genes and the world of phenotypes results in greater flexibility of the system architecture, where it is possible to adjust the degree of coexistence. It also allows for easy distribution of parallel evolutionary processes. The state of simulation is considered as the state of the gene pool only. When saving the evolutionary snapshot, the world is disregarded (the environment and all living entities represent an enormous amount of data), and only the gene pool is saved.

The creatures are revived based on the genotypes from the gene pool. Selection may be random or proportional to the weighted sum of the predefined criteria (if directed evolution is to be used). The directed model means that evolution is considered as the process of optimization, as it is in genetic algorithms (Goldberg 1989) and many artificial life simulations (Sims 1994). It may be assumed in simplified reality, that reproduction is generally proportional to the life span of individuals (and living requires energy acquisition). Then it is possible to mimic some aspects of open-ended evolution with the 'directed' model of evolution, when the fitness criterion (and selection) is directly connected with the survival and reproduction abilities: the life span criterion can be used as the estimate of reproduction and fitness.

Physical Simulation

Framsticks simulates a three-dimensional world and creatures. We decided to use such a sophisticated environment for evolution, expecting that a range of complex, various stimuli affecting organisms will be the origin of dynamic development and emergence of interesting behaviors. All kinds of interaction between physical objects are considered: static and dynamic friction, damping, action and reaction forces, energy losses after deformations, gravitation, and uplift pressure – buoyancy (in a water environment).

One should note that there is always a tradeoff between simulation accuracy and simulation time. We need a fast simulation to perform evolution, on the other hand the system should be as realistic (detailed) as possible to produce realistic (complex) behaviors. As we expect emergence of more and more sophisticated phenomena, the evolution has to be longer and the simulation must be less accurate, but faster.

Currently, in order to make the simulation fast and due to the computational complexity, some aspects were discarded: collisions between parts of an organism itself and the movement of a water medium were both ignored. Including these in the simulation would make it more realistic, but would not introduce significant, qualitatively new phenomena. Meaningful interactions were considered more important than very realistic, but too slow simulation (Maes 1995).

The basic element is a stick made of two flexibly joined particles (finite element method is used for simulation). Sticks have specific properties: biological (muscle strength, stamina, energetic: assimilation, ingestion, and initial energy level), physical (length, weight, friction), and concerning stick joints (rotation, twist, curvedness). Biological properties are mutually exclusive: an increase in one property of a stick results in a decrease of the rest, which is the price of specialization.

Muscles are placed on stick joints. There are two kinds of muscles: bending and rotating. Positive and negative changes of muscle control signal make the sticks move in either direction – it is analogous to the natural systems of muscles, with flexors and extensors. The strength of a muscle determines its effective ability of movement and speed (acceleration). A stronger muscle consumes more energy during its work.

Framsticks have currently three kinds of *receptors* (*senses*): those for orientation in space (equilibrium sense, gyroscope), detection of energy (smell) and detection of physical contact (touch). A sample framstick equipped with these elements is shown on figure 1.

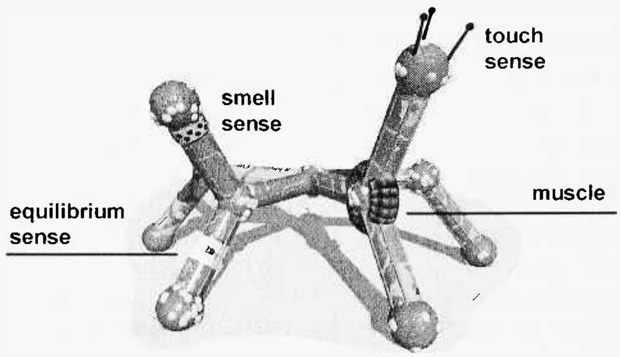

Fig. 1. Receptors (equilibrium, touch, smell) and effectors (muscles) in Framsticks

Neural Network

In order to expand the expression abilities of neural networks, the standard artificial neurons were enriched with a few additional parameters which could alter the behavior of each neuron independently. Since Nature was our inspiration, we did not introduce sophisticated and unnatural processing units – as in (Sims 1994) – because it is possible to construct complex modules (integrating, differentiating, summing, subtracting, and generators with different shapes) from simple neurons.

The additional neuron parameters (*force*, *inertia* and *sigmoid*) are under genetic control. They affect the way neurons work: *Force* and *inertia* influence changes of the inner neuron state (speed and tendency of changes, respectively). The *sigmoid* coefficient modifies the output function. Details and sample neuronal runs can be found at (Komosiński and Ulatowski 1997).

The neural network can have any topology and complexity. Neurons can be connected with each other in any way (some may be unconnected). Inputs can be connected to outputs of another neurons, constant values, or senses, while outputs can be connected to inputs of another neurons, or to effectors (muscles).

Genotype Representations

Various genotype encodings are proposed in order to study the processes of evolution in different setups, and to test their suitability for open-ended evolution. All of the representations combine "body" and "brain" in the same genotype, so that both morphology (body made of sticks) and control (brain made of neurons) evolve simultaneously. The simplest one is the basic direct encoding, which is least restrictive and fully exploits the capabilities of the world simulator (the genotype simply lists all the parts and details of an organism). Other encodings are described below.

Recurrent Direct Encoding

In this representation, the genotype describes all the parts of the corresponding phenotype. Small changes in the genotype cause small changes in the resulting creature. Stick phenotypic properties (see "Physical Simulation" in section III) are represented locally, but propagate through a creature's structure (with decreasing power). That means that most of the properties (and neural network connections) are maintained when a part of a genotype is moved to another place. Control elements (neurons, receptors) are associated with the elements under their control (muscles, sticks). Only tree-like structures can be represented (no cycles allowed).

While body is made of sticks, brain is made of neurons with their connections described *relatively*. Such a way of describing connections lets sub-networks come through the crossover operation: the whole set of neurons (a module) can be moved to another place in the genotype (and in the creature), possibly with limbs, and can still be operational.

The operations of mutation, crossing over and repair (to validate minor representation errors) are introduced.

Developmental Encoding

In nature, the genetic code of complex organisms does not encode their body layout directly, but rather their *process of development*. There are several hypothesized benefits of a developmental encoding in an evolutionary system (natural or artificial): higher evolvability, higher adaptation to environmental effects, compactness (Rotaru-Varga 1999). Based on these theoretical assumptions, we are developing a genotype encoding which is development-oriented, similar to encoding applied for evolving neural networks (Gruau 1996). An interesting merit of developmental encoding is that it can incorporate *symmetry* and *modularity*, features commonly found in natural systems, yet difficult to formalize.

The developmental encoding is similar to the recurrent direct one, but codes are interpreted as commands by cells (sticks, neurons, etc.). Cells can change their parameters, and divide. Each cell maintains its own pointer to the current command in the genetic code. After division, cells can execute different codes, and thus differentiate themselves. The final body (phenotype) is the result of a development process: it starts with an undifferentiated ancestor cell, and ends with a collection of interconnected differentiated cells (sticks and neurons). Codes affecting stick and neuron parameters are identical to the representation in the direct encoding. Currently the developmental process is not a part of the simulated world, but it is implemented as an external process during genotype decoding. However, a developmental process inside the simulation would add a new set of interactions between the environment and genes, through the environment affecting the development of a creature.

Other Encodings

To test various representations of creatures (morphologies and brains), and the influence of the representations on the evolutionary process, other encodings are also proposed.

In the "direct similarity development" encoding, all the parts of a creature are described in the genotype, but each part is a separate object. Each object has some multi-dimensional "links", which are connected together during the embryogeny process in the way that maximizes their similarity. Thus the final creature is developed.

In the "implicit embryogeny development" encoding, the genotype encodes a set of rules which are used during embryogeny to develop a creature. The rules concern spatial emission of some "chemicals", which affect growth of a creature and future activation of rules themselves.

Further experiments with all the encodings and their comparison are yet to be done. Another suggestions and ideas of representations can also be relatively easily incorporated in Framsticks.

Discussion of Results and Future Work

Our first experiments concerned the study of locomotion and orientation, so the fitness was defined as speed (on the ground or in water). Recurrent direct encoding was used. Many walking and swimming species evolved during evolutionary runs, and we were able to see the evolution of ideas of "how to move" (Komosiński and Ulatowski 1997).

In one evolutionary run, a limb of an efficiently-moving creature was doubled while crossing over, and after some further evolution the organism was able to move with two limbs – one for pushing back and one for pulling. We also noticed a case when a limb was simultaneously bent and rotated, which was a more effective method of pushing against the ground. In one case a neuron used its saturation to produce delayed signals, which is a kind of a simple "short-term memory". More sophisticated creatures which

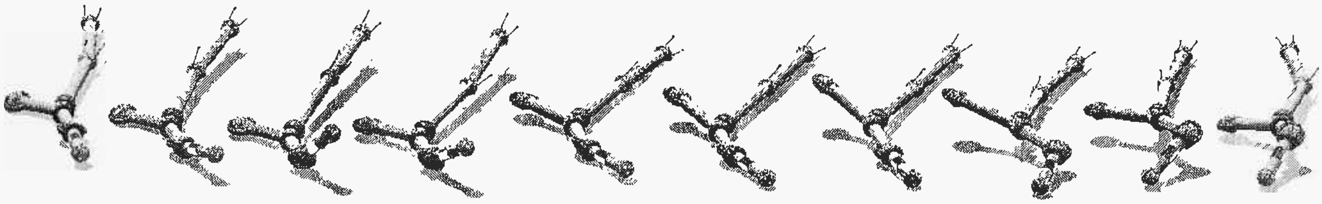

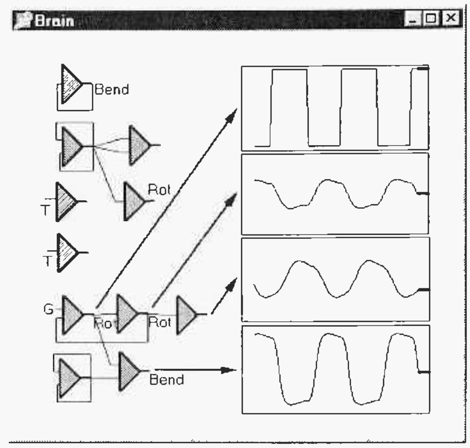

Fig. 2. A simple walking framstick. Two touch receptors (letter T), one gyroscope (G). Note the redundancy and random character of brain (six top-most neurons, including all touch sensors, do not influence the movement: touch neurons are unconnected, and the two muscle-controlling neurons are passive). Three (out of five) muscles are working; oscillations are produced only by the equilibrium sensor (G).

were evolved could not be easily examined because of their high complexity. A simple walking example is shown on figure 2. Even though most of body and brain designs seem to be redundant and not optimal, manual removing of some parts often makes individuals less fit. This happens due to various implicit dependencies (like feedbacks, etc.), which are overlooked or their influence is too complex to predict.

The experiments concerning evolutionary improvement of human predesigned structures have also been conducted. An example is the successful evolution of control (the neural network) for a hand-designed morphology in order to obtain a creature that lives long. Therefore, creatures were required to find energy sources and ingest them.

Some of the creatures display realistic behaviors even though they were not intended to imitate the real animals, as it was investigated in (Ijspeert 1999, Cruse et al. 1998). An example is a "salamander" (14 sticks, 16 neurons), which, after directed evolution of the neural network with fitness defined as speed, walks in a realistic way (see fig. 3). Another such example is swimming creatures; a swimming individual is studied in (Komosinski 2000).

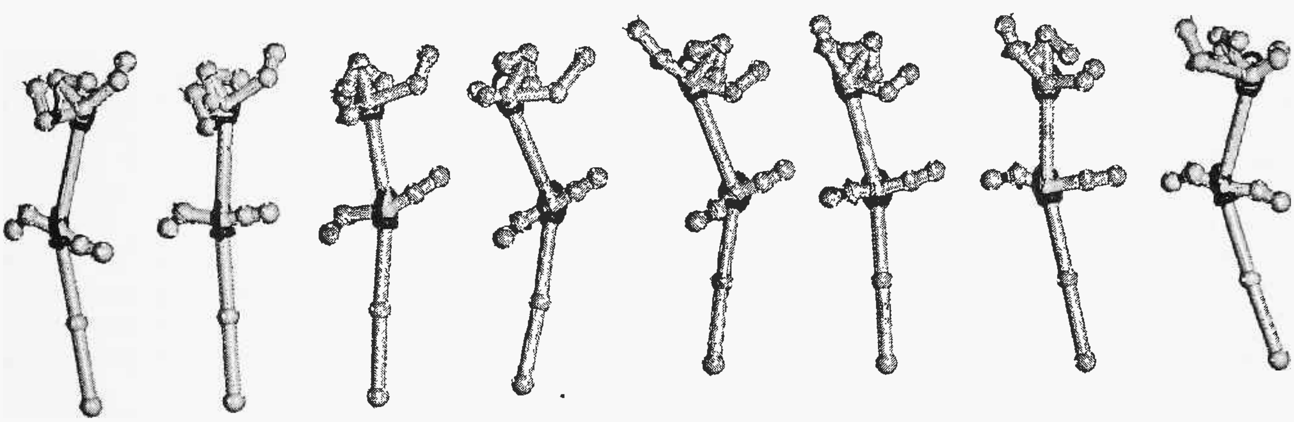

Fig. 3. A salamander-like, predesigned framstick. After evolution of brain walks in a realistic way.

A serious problem preventing open-endedness is when evolution is trapped in a local fitness maximum region. We attacked this problem from several sides: more evolvable genetic representations can induce a fitness landscape with less local peaks, and coevolution can dynamically change the fitness landscape.

When the survival of a creature does not depend on itself and the environment only, but also on interaction with other creatures (possibly from other species), the fitness landscape becomes dynamic (by coevolution). Effectively, a local maximum is 'lowered' by species trying to exploit it. Coevolutionary dynamics had been analyzed extensively in prey-predator (2 species) setups. The importance of coevolution is captured, among others, by (Steels 1997):

> Whereas evolution in itself causes an equilibrium to be reached, co-evolution causes a self-enforcing spiral towards greater complexity.

Framsticks has the potential for multi-species coevolution. Creatures can interact in various ways (push, eat, kill, etc.), and species can form and evolve simultaneously. However, a condition is that the number of simultaneously evolved creatures is high, which results in slow simulation.

We plan to improve the system architecture, and make the selection process more natural. A process of selection and reproduction which is external to the simulation does not resemble the natural situation. In a natural (or 'implicit') selection setup, creatures can trigger reproduction themselves when certain conditions are met (regarding energy, age). Thus the actual fact of survival is used instead of the probability of survival – numerical fitness (Cliff, Harvey and Husbands 1993). While natural selection makes it more difficult to 'breed' creatures to solve specific subtasks, it has better potential for maintaining perpetual evolution.

Our future work will also concern defining a similarity function on the phenotype space (to allow automatic analysis of populations and species, and to allow artificial speciation in directed evolution setup). We will compare various genotypic encodings and their influence on efficiency of directed and open-ended evolution. More receptors may be introduced. If the increase in computing power permits, we will use a more accurate (but slower) physical simulator, so that the evolved creatures will behave in a more realistic way.

Creating complex evolutionary systems, like Framsticks, raises the problem of analysis of the achieved results. After long efforts to *synthesize* artificial, realistic, open-ended ecosystems, we face the problem of *analyzing* something that was meant to become complex. Analysis of such results may rather be qualitative and behavioral, as simple quantitative measures are insufficient to capture the *way* and *method* of evolved successful individuals. Investigation of sophisticated creatures requires support of human-friendly tools and intelligent automatic systems, or human experts.

Thus we decided to develop a possibility of worldwide participation (Komosiński and Ulatowski 1997) in the process of setting up experiments, their distributed execution, and analysis of results. We encourage the ALife community to take part in this collective enterprise.

The achieved results are promising, and we hope to contribute to a large-scale, open-ended evolutionary experiments and analysis. We are looking forward to exploit full functionality of the simulator and learn about the results of such experiments.

Acknowledgements

Maciej Komosiński wishes to thank KBN, Polish State Committee for Scientific Research, for the research grant supporting this work.

References

Adami, C., and Brown, C. T. 1994. Evolutionary learning in the 2D artificial life system "Avida". In Rodney A. Brooks and Pattie Maes (eds.), Artificial Life IV: MIT Press/Bradford Books.

Balkenius, Ch. 1995. Natural Intelligence in Artificial Creatures, Ph. D. Thesis, Lund University Cognitive Studies.

Bajcsy, R. 1998. Active Perception. In Proceedings of the IEEE (76) 8, 996-1005.

Channon, A. D., and Damper, R. I. 1998. Perpetuating evolutionary emergence. In From Animals to Animats 5, Proceedings of the Fifth International Conference on Simulation of Adaptive Behavior. Rolf Pfeifer, Bruce Blumberg, Jean-Arcady Meyer, and Stewart W. Wilson (eds): MIT Press.

Channon, A. D., and Damper, R. I. 1998. Evolving Novel Behaviors via Natural Selection. In Artificial Life VI, Proceedings of the Sixth International Conference on Artificial Life. Christoph Adami, Richard K. Belew, Hiroaki Kitano, and Charles E. Taylor (eds): MIT Press.

Cliff, D., Harvey, I., and Husbands, P. 1993. Incremental evolution of neural network architectures for adaptive behaviour. In M. Verleysen (ed.), Proceedings of the First European Symposium on Artificial Neural Networks, 39-44. Brussels, Belgium: D-facto/Editions Quorum.

Cruse, H., Kindermann, T., Schumm, M., Dean, J., and Schmitz, J. 1998. Walknet – a biologically inspired network to control six-legged walking. In *Neural Networks* 11 1435-1447: Elsevier Science Ltd.

Goldberg, D. E. 1989. *Genetic Algorithms in Search, Optimization, and Machine Learning*: Addison-Wesley Publishing Co.

Gruau, F. 1995. Modular Genetic Neural Networks for 6-Legged Locomotion. In Artificial Evolution European Conference, Lecture Notes in Computing Science, 201-219. Alliot, J.-M., Lutton, E., Ronald, E., Schoenauer, M., Snyers, D., eds: Springer, 1996.

Hofstadter, D. R. 1979. *Goedel, Escher, Bach: an Eternal Golden Braid.* NY: Basic Books.

Ijspeert, A. J. 1999. Synthetic Approaches to Neurobiology: Review and Case Study in the Control of Anguiliform Locomotion. In Proceedings of 5th European Conference on Artificial Life, 195-204. September 13-17, 1999, Lausanne, Switzerland: Springer-Verlag.

Komosiński, M., and Ulatowski, Sz. 1997. Framsticks Internet site, documentation, and sample experiments, http://www.frams.poznan.pl/

Komosiński, M., and Ulatowski, Sz. 1999. Framsticks: towards a simulation of a nature-like world, creatures and evolution. In Proceedings of 5th European Conference on Artificial Life, 261-265. September 13-17, 1999, Lausanne, Switzerland: Springer-Verlag.

Komosiński, M. 2000. The World of Framsticks: Simulation, Evolution, Interaction. In Proceedings of 2nd International Conference on Virtual Worlds, Paris: Springer.

Maes, P. 1995. Artificial Life meets Entertainment: Interacting with Lifelike Autonomous Agents. In Special Issue on New Horizons of Commercial and Industrial AI 38, 11 108-114. *Communications of the ACM*: ACM Press.

Rotaru-Varga, A. 1999. Modularity in evolved artificial neural networks. In Proceedings of the 5th European Conference on Artificial Life, Lausanne, 1999. D. Floreano, J.-D. Nicoud, F. Mondada, eds: Springer.

Sims, K. 1994. Evolving 3D Morphology and Behaviour by Competition. In R. Brooks, P. Maes (eds.), Artificial Life IV Proceedings, 28-39: MIT Press.

Steels, L. 1997. Synthesising the origins of language and meaning using co-evolution, self-organisation and level formation. In: Hurford, J., C. Knight, and M. Studdert-Kennedy (ed.), Evolution of Human Language: Edinburgh Univ. Press.