
Population dynamics and emerging mental features in AEGIS

A.E. Eiben

Leiden University
The Netherlands
gusz@cs.leidenuniv.nl

D. Elia

Antares Informatisering B.V.
The Netherlands
delia@antares.nl

J.I. van Hemert

Leiden University
The Netherlands
jvhemert@cs.leidenuniv.nl

Abstract

We describe an empirical investigation within an artificial world, AEGIS, where a population of animals and plants is evolving. We compare different system setups in search of an ‘ideal’ world that allows a constantly high number of inhabitants for a long period of time. We observe that high responsiveness at individual level (speed of movement) or population level (high fertility) are ‘ideal’. Furthermore, we investigate the emergence of the so-called mental features of animals determining their social, consumptional and aggressive behavior. The tests show that being socially oriented is generally advantageous, while aggressive behavior only emerges under specific circumstances.

1 Introduction

The study we report here is concerned with the population dynamics and the emergence of mental features¹ under different conditions in an A-Life environment. Technically speaking we have two objectives:

1. To investigate which environmental conditions (i.e. physical and biological rules) are ‘ideal’, in the sense that the ecosystem is able to maintain a constantly high number of inhabitants for a long period of time.
2. To study the emerging dominant patterns in the mental characteristics of artificial animals.

¹Mental features are the ones that directly determine behavior of individuals.

Our artificial world is called Artificial Environment with Genetic Inheritance Simulation (AEGIS)². It is a simple model of a natural ecological system for simulation and experimentation purposes, no optimization problems of any kind are solved like in (Hillis, 1991). This system embodies an artificial environment with two kinds of inhabitants: *animals* and *plants*. The animals are genetic which means that their physical and mental properties are determined by their chromosomes, and genetic operators, such as crossover and mutation, are used on the chromosomes to produce offspring. In (Koza, 1991), (Rasmussen et al., 1991), (Jefferson et al., 1991), (Koza, 1994), (Ray, 1991), and (Skipper, 1991), the genetic information is based on programmable matter like executable code. The genetic information of the animals in our environment is based on DNA and RNA like structures (Schuster, 1991). The plants are not genetic and are primarily incorporated in the system to serve as food for the animals. Unlike (Taylor et al., 1988), AEGIS has biological knowledge already built in. The animals are able to move through the artificial environment; they are not stationary like cellular automata ((Tamayo and Hartman, 1988), (Boerlijst and Hogeweg, 1991), (Kauffman and Johnsen, 1991)). In Section 2 we will give a description of AEGIS.

Certain laws of physics and biology within AEGIS are incorporated by general rules (describing for instance how animals move) and by parameters used by these rules. We study the effects of different parameter settings and the effect of a certain parameter is observed at two levels. For our first goal we look at the population dynamics, starvation and birth figures for instance, or periodicity in the system. For our second goal we look at the mental characteristics of the animals at the end of the run. If a certain characteristic, e.g. much interest in food, is commonly present, than

²AEGIS (written in C++/Motif) can be downloaded from <http://www.wi.leidenuniv.nl/~jvhemert/aegis/>

we see it as an essential feature for survival in the given environment. The experiments are described in Section 4.

Finally, in Section 5 we summarize our observations, draw conclusions and sketch interesting issues for further research. This paper leaves out much of the implementation details of the AEGIS world. For a more detailed description we refer to (Elia, 1997a) and (Elia, 1997b). Similar work has been done in (Bedau and Packard, 1991), (Devine and Paton, 1997), (Coderre, 1988), (Johnson, 1994) and (Skipper, 1991).

2 AEGIS

The artificial world can be seen as a two dimensional grid wrapped around at the edges (Figure 1). Each grid position corresponds with an area which can contain one animal and one plant.

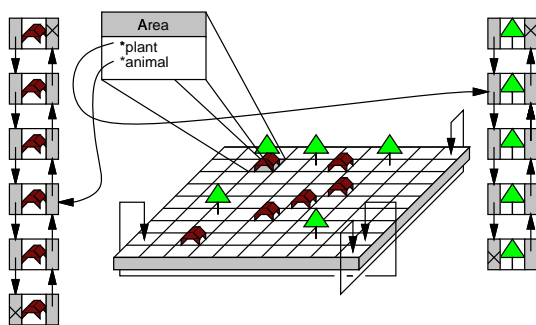


Figure 1: The Artificial World structure.

The world evolves with discrete time steps, called cycles. In one cycle all animals are given the opportunity to perform a number of actions. The number of actions an animal can perform depends on its physical attributes (more muscles imply more actions per cycle, while having a thick skin implies fewer actions). After each cycle, the number of actions animals can perform in the next cycle is updated, as well as the plant attributes. Animals are born with a certain maximum life time. After each cycle the remaining life time is updated by reducing the maximum life by one. (The remaining life time is also reduced by receiving a hit in a fight, see Section 3.2.)

Artificial Plants (APs) are necessary for animals as food to restore their energy reserves. An AP has a certain fixed position in the environment and an energy level which equals its nutritious value for an animal. APs have three features: reproduction, growth and mortality. At the end of a cycle APs grow. If an AP's energy level exceeds the multiplication threshold

it will try to create a new AP in one of its four neighboring areas. When this is not possible and if the energy level exceeds the seed launch threshold, it will randomly try to create a new AP somewhere in the world, killing itself in the process. An AP also dies when it is completely eaten by an animal, which results in an energy level of zero.

The main subjects in this study are the Artificial Animals (AAs). AAs mimic important features of real living organisms: mobility, reproduction, competition, consumption and mortality. In the system, the plants are the energy providers and the animals are the energy consumers. The animals acquire energy by eating plants and lose energy by performing actions. Possible actions are: moving, eating, fighting and reproducing. An AA can observe (part of) its environment. An evaluation function, influenced by the mentality of the given AA, decides which direction of movement is most advantageous. AAs have genetic information (genotype), which determines their physical (muscle mass and skin thickness) and mental attributes (interest levels in social, aggressive and food consuming behavior). The genetic material is encoded in bit-strings which are subject to genetic operators. One point crossover is applied to two parents and delivers two children. Children are created by cutting the parents into two pieces at a randomly chosen crossover point and exchanging the genetic material after this point. Mutation is applied to one animal and it creates one new animal by changing the value of a bit at a randomly selected position. Every time an animal performs an action there is a chance of 0.01 this animal gets mutated. It is not possible to mate at just any given time. Animals have to wait for a so-called recovery period after having created offspring before they are able to reproduce again.




Figure 2: An illustration of a genotype of an AA, with physical (gender, muscle, skin) and behavioral (social, consumptional, aggressive) parts.

3 Animal behavior

We distinguish three types of behavior: social, consumptional, and aggressive (attacking) behavior. The mental bits of a genome are used by probabilistic rules in the decision mechanism shown in Figure 3. If any of

the behavioral attributes has a high value, the AA will have the tendency to behave in that particular way. Nevertheless, by the probabilistic nature of the rules the choice to perform a certain action is not deterministic.

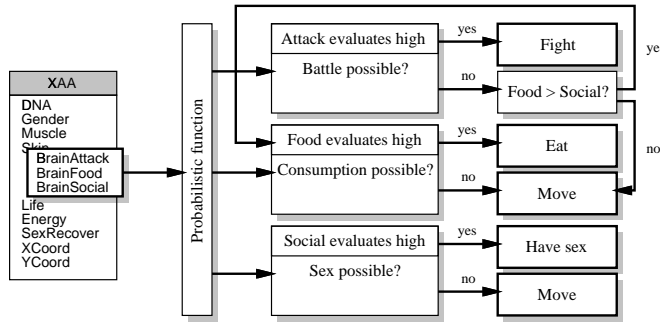


Figure 3: Decision mechanism of an AA.

AAs can interact with each other in various ways. The type of interaction is based on the difference in physical attributes (skin thickness and muscle strength). When an AA meets another AA which has physical attributes that are outside the tolerance bound (here set at 12.5%), they consider each other as rivals making aggressive behavior possible. AAs with similar attributes have the possibility to ‘socialize’. Because the physical attributes are also subject to evolution, an implicit speciation process is taking place during a run. These aspects fall outside the scope of this paper, here we only mention that thin skin and high muscle mass relatively increase speed (the number of actions per cycle) and lower the sex recovery period of an AA, (Eiben et al., 1998).

3.1 Social behavior

The tendency to show social activity is expressed by an AA’s social mental bits. The influence of the social interest is twofold. First, an AA with social interest will have the tendency to move towards other AAs which are physically similar (within the tolerance bound). Second, when an AA has a high social interest, it will try to have more sex. Two AAs can only mate if they are on neighboring areas and

- the gender attributes differ,
- the physical attributes are within the physical tolerance bounds (here 12.5%),
- the mental attributes are within the mental tolerance bounds (here 25.0%).

After mating, AAs are not able to mate again for a so-called a recovery period. This period is measured by the numbers of actions performed and not by the number of cycles elapsed. Thus, speedy AAs recover in fewer cycles.

3.2 Aggressive behavior

Two AAs residing in neighboring areas and being outside each others physical tolerance bound can have a fight or decide to run off. An AA will fight if it has an aggressive attitude and if it has enough energy. A fight consists of a series of hits. Each hit reduces the energy level of the attacker by **fight-cost** units and also lowers the energy level of the victim by two times **fight-cost** units. Another effect of a hit is that it reduces the remaining life time of the victim by a number of cycles. This number is determined by the muscle mass of the attacker (high muscle mass can cause more damage) and the skin thickness of the victim (thick skin animals are more resistant).

3.3 Consumptional behavior

Consumptional attitude has its influence on the decisions and movements of an AA. A high interest in food consumption generates a tendency to move towards areas containing plants by influencing the evaluation function that determines how attractive given areas are. Additionally, when an AA has less than half of its maximum energy reserves, the attractiveness of areas containing food is increased ten times to prevent starvation. When an AA eats, it extracts energy from the plant on the current area. The amount of energy it extracts is maximal in the sense that it either refills the AAs reserves to maximum or exhausts all the plant’s energy.

4 Experiments

We performed initial experiments in order to adjust the system parameters such that the eco-system as a whole is viable, i.e. the amount of animals and plants does not becomes zero. A configuration satisfying this requirement is taken as the reference parameter set for AEGIS. Parameters are divided into fixed parameters, which are kept fixed throughout all experiments (see Table 1) and variable parameters (shown in Table 2).

The effects of the variable parameters are the subjects of the current study. The first three parameters have effect on the animals’ metabolism: the amount of energy units it takes to perform a certain action. The last parameter has influence on the fertility, by de-

Table 1: Values of fixed system parameters.

Animal status specification	
maximum life time (nr. of actions)	200
maximum energy reserve (energy)	150
Initial population	
Initial Animal population size	1000
Initial Plant population size	3000
Plant characteristics	
Plant energy at ‘birth’ (energy)	75
Plant growth (energy/cycle)	20
Plant multiplication threshold (energy)	125
Plant seed launch threshold (energy)	140

terminating the frequency of reproduction and thereby the amount of offspring an animal can create in a lifetime. The reference parameter set has been used to depict the ‘standard’ evolution used as a yardstick to compare the other parameter sets with.

Table 2: Values of variable system parameters in the reference parameter set.

Action costs	
Movement cost (energy/movement)	5
Sex cost (energy/mating activity)	10
Fight cost (energy/hit)	5
Sex recovery (actions)	20

A system run is monitored in different ways. Basically, we consider population dynamics/statistics, and the emergence of animal attitudes. Population dynamics/statistics are monitored by recording data on the number of animals, plants, fights, children born and animals died during the run. This kind of data (gained by five independent runs for each setup) is summarized in Table 4, where the first row contains the reference parameter set. In Figure 4 we display one run of the system by plotting the curves on the number of animals, plants, and fights. Note that such curves cannot be averaged, if two curves behave the same way, but are relatively out of phase, they can nullify each other’s waves. Thus, instead of showing averaged curves, a representative run is shown. Notice the influence the animals and plants have on each other. In the first stage of the run (cycles 1 through 500) the system is in some kind of transition, or ‘warming-up’ phase. After that, the curves show low frequency repetitions in a predator-prey style. Overpopulation causes a decrease in plants. This will cause an increase in deaths, which means that the plant life can fully restore eventually.

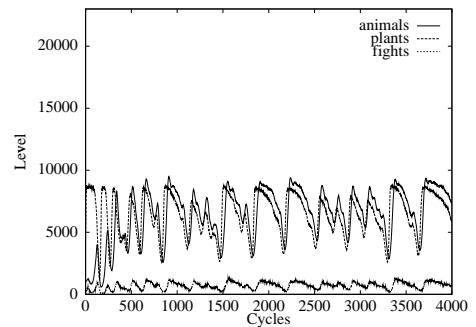


Figure 4: Population dynamics under the reference parameter set (one representative run shown).

Besides population dynamics/statistics we also look at the development of animal characteristics under the circumstances represented by a certain parameter set. We distinguish physical attributes (skin thickness and muscle mass) and mental attributes (interest in food consumption, socializing and aggressive actions). By the lack of space we omit the treatment of physical attributes, and concentrate on mental features in this paper. A concise overview of results can be found in (Eiben et al., 1998). Since all animal features are initialized randomly, data in the beginning of a run show almost equally divided values for each attribute. Non-uniform distributions at termination indicate that under the selective pressure of the environment certain characteristics disappear, thus are considered as disadvantageous, while others prevail. These are seen as successful survival strategies in the setup.

The mental features are given by histograms showing the end situation (after 4000 generations) in Figure 5. The three histograms depict the tendency for aggression (left), consumption (middle), and social interest (right) in the population. In each histogram the left-hand side belongs to lower values concerning the given feature. That is, if a histogram is skewed to the left then the AAs are not much interested the given kind of activity (fights, eating, or sex). Figure 5 shows that the dominating attitudes are not aggressive (cautious/neutral) and socially interested. A consumptional attitude does not seem to be that important since the values vary between very low and very high.

4.1 Comparing different parameter sets

To examine the effect of the parameters move cost, sex cost, fight cost, and sex recovery have on the artificial ecosystem, we systematically change them. The system parameters undergo isolated adjustments, i.e. only




Figure 5: Histograms of mental animal attributes at termination under the reference parameter set: aggression (left), consumption (middle), and social interest (right).

one parameter is changed at a time. For each of the above parameters we experiment with a low and a high value as compared to the reference parameter set. Table 3 shows the particular values that have been used.

Table 3: The variable system parameters for all setups (setup 1 is the reference parameter set).

nr.	name	move cost	sex cost	fight cost	sex recov
1	reference set	5	10	5	20
2	low fight cost	5	10	1	20
3	high fight cost	5	10	20	20
4	low move cost	1	10	5	20
5	high move cost	20	10	5	20
6	low sex cost	5	1	5	20
7	high sex cost	5	40	5	20
8	short sex recov	5	10	5	1
9	long sex recov	5	10	5	30

In the sequel we summarize the results of changing the parameter values, based on five independent runs for each setup. Recall, however, that the system curves cannot be averaged. For the population dynamics a representative run is shown in all cases.

Table 4: Average number of animals, plants, fights, offspring born, and deaths over five independent runs for all setups.

	animals	plants	fights	offspring	deaths
1	6623.8	6741.1	861.4	155.2	73.1
2	7109.4	6838.9	1356.9	165.2	58.8
3	4226.7	6103.5	344.1	111.0	98.5
4	9379.4	7897.4	3606.2	322.7	41.5
5	1138.2	7060.7	48.0	53.0	55.9
6	7006.7	6590.3	797.0	178.5	70.8
7	5663.6	6452.0	523.8	129.5	69.3
8	9165.4	7815.1	1703.8	196.4	75.7
9	1388.8	7223.1	266.1	49.9	34.2

4.1.1 The effect of fight cost

The most significant difference between the low and high fight cost setups is the frequency of oscillations in the number of animals and plants. High costs increase this frequency and, to a less extent, also the amplitude. It is interesting that aggressive attitude only emerges under high fight costs. This increased interest in fighting costs a lot of energy. The increased interest in food, cf. Figure 7, and the higher death figures, cf. Table 4, confirm this.




Figure 6: Population dynamics for parameter set 2: low fight cost (left) and parameter set 3: high fight cost (right).




Figure 7: Histograms of mental animal attributes at termination for low fight cost (above) and high fight cost (below): aggression (left), consumption (middle), and social interest (right).

4.1.2 The effect of move cost

Here again, we see that the setup with the high cost (in this case high move cost) leads to an increased frequency of oscillations in the number of animals and plants. The low move cost setup seems ideal: the ecosystem is able to reach a stable state and maintain animal and plant populations at a high level. Very surprising is the emergence of aggressive behavior in case of the low move cost setup. There are even more aggressive animals than under high fight costs (Figure 9) and the number of fights is also much higher, cf. Table 4. The increased move cost leads to an increase interest in food.




Figure 8: Population dynamics for parameter set 4: low move cost (left) and parameter set 5: high move cost (right).




Figure 9: Histograms of mental animal attributes at termination for low move cost (above) and high move cost (below): aggression (left), consumption (middle), and social interest (right).

4.1.3 The effect of sex cost

The difference between the low and high cost setup variants is the smallest so far in these experiments. Nevertheless, the high sex cost setup shows an increased frequency w.r.t. the one with low cost. The relatively small difference in the results in Table 4 and the emerging mental characteristics indicate that the sex cost is a robust parameter.




Figure 10: Population dynamics for parameter set 6: low sex cost (left) and parameter set 7: high sex cost (right).

4.1.4 The effect of sex recovery period

The difference between these two setups is dramatic. The short sex recovery setup qualifies as the ‘second best paradise’ (after low move cost), while having a long sex recovery period leads to a complete disaster:




Figure 11: Histograms of mental animal attributes at termination for low sex cost (above) and high sex cost (below): aggression (left), consumption (middle), and social interest (right).

the eco-system dies out after about 1100 cycles³. The high fertility implied by the short sex recovery period enables an almost stable state (Figure 12), where the population levels are close to those in the low move cost setup (see Table 4). There are also significant differences between the two ‘ideal’ setups: as opposed to parameter set 4, no aggressive attitude emerges under parameter set 8. Furthermore, in the low sex recovery setup the interest for food is significantly higher as shown in Figure 13.




Figure 12: Population dynamics for parameter set 8: short sex recovery (left) and parameter set 9: long sex recovery (right).




Figure 13: Histograms of mental animal attributes at termination for short sex recovery period (above) and for long sex recovery period (below): aggression (left), consumption (middle), and social interest (right).

³Note, that all statistics and data plots for this setup are based on the situation after 1000 cycles, as opposed to the usual 4000.

5 Conclusions

Our first research objective concerns the parameters of an ‘ideal’ world that allows a constantly high level of flora and fauna in AEGIS. We have found two of such worlds, belonging to parameter sets 4 (low move costs) and 8 (short sex recovery period). These worlds are not identical, since the curves on population dynamics (the left hand side of Figure 8, respectively the left hand side of Figure 12) differ, and so do the figures from Table 4. Nevertheless, they are clearly different from the other ones and similar to each other. These worlds share a common feature: the systems ability for fast response. Low move costs allow fast individuals (many movements per cycle on the two dimensional grid), while a short sex recovery period facilitates the population’s fast traversal of the imaginary space of all possible genotypes. In other words, the ideal world seems to be one that allows fast action at individual or at population level. Let us explicitly mention the world with low fertility (long sex recovery period). This parameter setup turns out to be lethal, since it leads to complete extinction.

It is interesting to note that worlds with high costs (for fight, move and sex) lead to more frequent oscillations in the populations compared to the low cost variants, see the left hand side and the right hand side of Figures 6, 8, 10, and 12. Different values for the sex cost parameter, however, caused relatively little differences in the statistics (Table 4) and the population dynamics (Figure 10), that is, for other parameters the low and high valued setups exhibit larger differences. The small effect of different values of the sex costs are even more apparent for the mental features, compare the histograms in Figure 11 that are almost identical. Thus, sex cost seems to be a robust parameter, in the sense that the system is not too sensitive to its specific values in the range we have tried.⁴

As for our second research objective, the emerging mental features, probably the most interesting observation is that under all circumstances the animals develop a high social interest, see the rightmost histograms in figures 7, 9, 11, and 13. With respect to the animals interest in food there were large differences in emerging mental patterns. In general it seems that environments with high cost involve a consumption oriented behavior. Nevertheless, interest in food under low sex costs is also high.

Aggressive attitude only emerged under setup 3 (high

⁴Note that the range of different values is the highest for this parameter: between 1–40, while for the others the values range between 1–20 or 1–30.

fight costs) and setup 4 (low move cost), in other cases the animals are mostly cautious/neutral. The general lack of aggressive attitude is interesting as fights seem to offer evolutionary benefits. Namely, fights can reduce the energy level and the remaining lifetime of ‘enemies’, that is of AAs that are consuming the same resources, but differ so much physically that are useless for mating and reproduction. Besides this possible benefit, fights also have an unattractive side: costs. Recall that an AA only initiates a fight if its energy reserves are high enough and each hit in a fight reduces the energy level of the attacker too. In the meanwhile, the damage to the victim is larger, it loses twice as much energy and also some of its remaining life time. All in all, fighting is not attractive if the costs are relatively high compared to the benefits. Thus, fighting might be worth if the costs are not high (plenty of energy is available) or the benefits (damage caused) are large. The two worlds where aggressive attitude emerged illustrate both options. Under setup 3 (high fight costs), the difference between the losses of the attacker and the victim are the largest, which amounts to relatively large benefits. Setup 4 (low move costs) embodies a ‘cheap’ world where movement does not cost much energy, hence the cost aspect does not serve as prevention against fighting.

Further research is directed to incorporate the notion of communication in the world of AEGIS and study its effect on the dynamics of the population. Details will be subject of future publications, let us only mention one remarkable result here. Without communication world 9 (long sex recovery) ended up in extinction of all AAs, cf. the left hand side of Figure 12. Nevertheless, the possibility to communicate, in particular to inform each other about the location of food, results in a viable ecosystem under the same conditions as shown in Figure 14.




Figure 14: Population dynamics for parameter set 9 (long sex recovery) with communication added.

References

- Bedau, M. A. and Packard, N. H. (1991). Measurement of evolutionary activity, teleology, and life. In Langdon, C. G., Taylor, C., Farmer, J. D., and Rasmussen, S., editors, *Artificial Life II*, volume X. Addison-Wesley.
- Boerlijst, M. and Hogeweg, P. (1991). Self-structuring and selection: Spiral waves as a substrate for prebiotic evolution. In Langdon, C. G., Taylor, C., Farmer, J. D., and Rasmussen, S., editors, *Artificial Life II*, volume X. Addison-Wesley.
- Coderre, B. (1988). Modeling behavior in petworld. In Langdon, C. G., editor, *Artificial Life*, volume VI. Addison-Wesley.
- Devine, P. and Paton, R. (1997). Herby: An evolutionary artificial ecology. Technical report, Department of Computer Science, University of Liverpool, PO Box 147, Liverpool L69 3BX, UK.
- Eiben, A., Elia, D., and Hemert, J. (1998). The effect of communication in artificial life systems. Technical Report TR-98-12, Leiden University. Also available at <http://www.wi.leidenuniv.nl/TechRep/1998/>.
- Elia, D. (1997a). *The AEGIS User Manual*. Leiden University.
- Elia, D. (1997b). Population dynamics with and without communication in AEGIS. Master's thesis, Department of Computer Science, Leiden University, P.O. Box 9512, 2300 RA Leiden, The Netherlands.
- Hilis, W. D. (1991). Co-evolving parasites improve simulated evolution as an optimization procedure. In Langdon, C. G., Taylor, C., Farmer, J. D., and Rasmussen, S., editors, *Artificial Life II*, volume X. Addison-Wesley.
- Jefferson, D., Collins, R., Cooper, C., Dyer, M., Flowers, M., Korf, R., Taylor, C., and Wang, A. (1991). Evolution as a theme in artificial life: The genesys/tracker system. In Langdon, C. G., Taylor, C., Farmer, J. D., and Rasmussen, S., editors, *Artificial Life II*, volume X. Addison-Wesley.
- Johnson, A. R. (1994). Evolution of a size-structured, predator-prey community. In Langdon, C. G., editor, *Artificial Life III*, volume XVII. Addison-Wesley.
- Kauffman, S. A. and Johnsen, S. (1991). Co-evolution on the edge of chaos: Coupled fitness landscapes, poised states, and co-evolutionary avalanches. In Langdon, C. G., Taylor, C., Farmer, J. D., and Rasmussen, S., editors, *Artificial Life II*, volume X. Addison-Wesley.
- Koza, J. R. (1991). Genetic evolution and co-evolution of computer programs. In Langdon, C. G., Taylor, C., Farmer, J. D., and Rasmussen, S., editors, *Artificial Life II*, volume X. Addison-Wesley.
- Koza, J. R. (1994). Artificial life: Spontaneous emergence of self-replicating and evolutionary self-improving computer programs. In Langdon, C. G., editor, *Artificial Life III*, volume XVII. Addison-Wesley.
- Rasmussen, S., Knudsen, C., and Feldberg, R. (1991). Dynamics of programmable matter. In Langdon, C. G., Taylor, C., Farmer, J. D., and Rasmussen, S., editors, *Artificial Life II*, volume X. Addison-Wesley.
- Ray, T. S. (1991). An approach to the synthesis of life. In Langdon, C. G., Taylor, C., Farmer, J. D., and Rasmussen, S., editors, *Artificial Life II*, volume X. Addison-Wesley.
- Schuster, P. (1991). Complex optimization in an artificial rna world. In Langdon, C. G., Taylor, C., Farmer, J. D., and Rasmussen, S., editors, *Artificial Life II*, volume X. Addison-Wesley.
- Skipper, J. (1991). The computer zoo – evolution in a box. In Varela, F. J. and Bourgine, P., editors, *Toward a Practice of Autonomous Systems: proceedings of the first European Conference on Artificial Life*, volume XVII. The MIT Press.
- Tamayo, P. and Hartman, H. (1988). Cellular automata, reaction-diffusion systems and the origin of life. In Langdon, C. G., editor, *Artificial Life*, volume VI. Addison-Wesley.
- Taylor, C. E., Jefferson, D. R., Turner, S. R., and Goldman, S. R. (1988). Ram: Artificial life for exploration of complex biological systems. In Langdon, C. G., editor, *Artificial Life*, volume VI. Addison-Wesley.