

DELIBERATE EVOLUTION AGENTS: COMPARING REPRODUCTION STRATEGIES

C.M. Jonker, A.P.G. de Kock, J. Meijer, B.J. Vermeulen
Vrije Universiteit Amsterdam, Department of Artificial Intelligence
De Boelelaan 1081a, 1081 HV Amsterdam, The Netherlands
Email: {jonker, jmeijer, bjvermeu}@cs.vu.nl
URL: <http://www.cs.vu.nl/~jonker>

Abstract – In this paper the difference between two reproduction strategies of deliberate evolution is studied: partner selection and environment evaluation. The experiments conducted indicate that environment evaluation outperforms partner selection when timing of reproduction is important. When timing is not important, partner selection performs better than environment evaluation.

1 INTRODUCTION

The adaptive power of evolution has been an inspiration to artificial intelligence for some decades, e.g., [3]. In traditional genetic algorithms partners come together by an invisible hand that selects the most “fit” agents to mate with each other. In nature evolution is a complex process in which the agents involved play a more active role [6].

Reading work reported in literature shows that, next to natural selection, the suitability of the environment and the attractiveness of possible partners influence natural evolution processes [2,7]. For example, Miller and Todd [8] explain that natural selection is rather noisy, due to a diverse range of possibilities of nature’s selection procedure. Mate selection makes it possible to put the ancestral pressures of nature back into the offspring, thus cancelling the noise. In nature these two aspects influence evolution simultaneously, which makes it hard to understand the effect of each aspect in separation. Research in Artificial Intelligence and Artificial Life enables the isolation of aspects.

The effect of environmental structure on evolutionary adaptation is, for example, studied in [5]. The environmental structure studied there is the way the food is distributed over space and time. Fletcher, Bedau and Zwick show that more detectable environmental structure leads to better adaptation. To reach this conclusion, they simulated a world in which agents have to jump to a nearby site in the environment to reach food. The agents cannot see the exact food amount on each spot so they have to estimate where they can find most food. Experiments were done with environments with different structures. Their experiments show that environments with most information result in the best adaptation.

The way sensory capabilities of agents influence their adaptation to a given environment is studied in [11, 13]. In both experiments agents have with no sensors at all, but with just an action probability list. The list consists of seven values representing probabilities of choosing a certain action (like move right, move left, and procreate). This is used in the

evolution of the agents, performing mutation on the list when an agent splits. It is shown that even without the capability of sensing, such agents are capable of adapting to specific environments. For instance, with certain density settings, agents adapted their movements to clock-wise movements.

The effect of sexual selection on evolutionary adaptation is studied by Ventrella [14, 15]. Ventrella simulated a population of organisms which evolve morphology and motor control for fluid locomotion. In this simulation there was no explicit fitness function for the individuals. Individuals selected their mates on the basis of physical properties of the other agents. In Ventrella’s experiments the fitness of an individual is not only determined by the capability of the agent to survive, but also by the capability to attract other agents. Properties that may be attractive to agents, are not necessarily properties that help the agent to survive. In experiments where agents were attracted to lack of movement, evolution of efficient energy use was inhibited.

In this paper two different strategies are studied in which agents procreate deliberately. These agents, called deliberate evolution agents (DEA), can (to some extent) determine when to mate and with whom. In the DEA experiments no invisible hand selects partners. The strategies resemble two aspects that influence procreation in nature. The first is the appropriateness of the animals’ environment. For example, most animals are born in spring when food starts growing again. The second criterion studied is the attractiveness of the potential partner. The peacock, for example, is an animal that goes to great lengths to attract partners.

The DEA project reported here set out to research which of these criteria is more important. The question is whether a species evolves better when individuals use the environment as criterion or when the partners’ attractiveness is used.

Section 2 describes experimental set-up of the DEA project. The experiments conducted and results thereof are presented in section 3. In section 4 the conclusions of the DEA project are given and compared to existing literature.

2 EXPERIMENTAL SETUP

As Wilson [16] pointed out, the emphasis of the project should not lie in designing the agents to fit the world, but on the analysis of the resulting behaviour of the agents. Before describing the different types of DEA agents, first the DEA world is specified.

2.1 The DEA world

The world used in the experiments is formed by a rectangle size thirty in square. At regular time intervals 15 units of food grow at random places. Several agents and several pieces of food can occupy the same spot at the same time. Agents have to find food in order to survive. When an agent dies, it just disappears. Time in the DEA world is discrete time steps called cycles. In each cycle all agents have the opportunity to perform one action. Actions are performed sequentially; youngest agents first, oldest last. Because the food-supply is limited, agents have to compete for survival. The fittest agents are able to produce more offspring so “good” genes have a bigger chance to survive than bad ones.

2.2 DEA Agents

The agents in the DEA world can eat, find food, find a partner, mate and sleep. All actions cost energy and time as in the table below (Note that eating is the only action that provides energy, in stead of costing

TABLE 1 ACTIONS LIST

Action	Time consumption	Energy consumption
Eat	1	-20
Find food	4	2
Find partner	4	3
Mate	1	30
Sleep	8	1

Eating is possible when food is at the same spot as the agent looking for it. When an agent eats, the number of food units on the spot decreases by one. Eating ends the action of finding food. The energy level of the agent increases if it consumes a unit of food. When an agent decides to find food

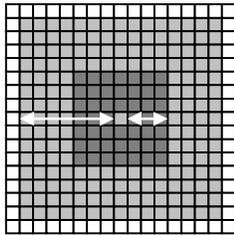


FIGURE 1
VISUAL MATING RANGE

The left arrow shows the visual range (by default 7), the right arrow shows the mating range (by default 3) The area that the agent can observe is light grey. To be able to mate, the desired partner has to be in the dark grey area.

and no food is available at the agent’s location, the agent moves into the direction of the nearest food unit. If no food is present within the visual range (see Figure 1) of the agent, the agent does not move at all.

When looking for a partner, an agent moves towards the nearest agent. The action ends when the agent is at the same spot as some other agent. As with finding food, if there is no other agent within visual range, the agent does not move. Mating can only occur when another agent is within mating range (see Figure 1), both agents have enough energy and the number of agents in the neighbourhood does not exceed a preset maximum. It is not necessary that both agents want to mate but mating costs energy for both parents. An agent has a maximum age of (by default) 200 cycles. The agent dies when its energy level is below zero or when it has reached the

maximum age. The energy of an agent that has just been born is 20 units.

2.2.1 The agents’ genes

The agent’s actions are listed in an action array, which is part of its DNA. Each array consists of a sequence of 10 actions. All actions are performed sequentially; when the agent finished the last action in the list, it starts over again. This means that an individual agent is not able to adapt its behaviour during its lifetime. Adaptation can only be achieved by evolution. Agents also have weights used to decide whether to mate or not. Weights are placed on a part of the chromosome called the property-array.

2.2.2 Reproduction

In order to decide whether to reproduce or not, the agent uses its weights. The agent observes three properties of the environment or other agents. Each of these observations has a corresponding weight, dictated by the agent’s genes. When the evaluation exceeds a threshold (and all other requirements are met), the agent mates. The agent evaluates the following:

$$\sum_{i=1}^3 o_i \cdot w_i > t_g$$

Where o_i stands for observable i , and w_i stands for the weight corresponding o_i . The global threshold t_g is set to 50. All weights are floats between -10 and 10 .

When an agent decides to mate, the genetic information of both parents is combined to produce the genes of the baby agent. One point crossover is applied to both the action array and the property array. This means that the chromosome is split at a random point. The baby receives one part of its genes from one parent, and the rest from the other.

When this crossover is performed, each gene mutates with a probability of 0.08. This mutation is done by random re-initialisation. This means that the gene gets a random new value.

2.2.3 Agent types

Four agent types were used. The main types were partner selection (PS) and environment evaluation agents (EE). These are both divided in static and relative agents. The agents differ in the observations on which they base the decision to mate or not.

PS relative agent

$$o_1 = age_i - age_{average}$$

$$o_2 = c_i - c_{average}$$

$$o_3 = e_i - e_{average}$$

Where age_i stands for the age of agent i , and $age_{average}$ for the average age of the population, similarly c stands for the number of children, and e for the energy level.

PS static agent

$$o_1 = age_i$$

$$o_2 = c_i$$

$$o_3 = e_i$$

EE relative agent

$$o_1 = a_{vr} - \frac{r_v}{s} \cdot a_{all}$$

$$o_2 = a_{mr} - \frac{r_m}{s} \cdot a_{all}$$

$$o_3 = f_{vr} - \frac{r_v}{s} \cdot f_{all}$$

Where a stands for the number of agents, in respectively virtual range, mating range or the whole world. Virtual range area and mating range area are denoted by r_v and r_m respectively. The area size of the world is indicated by s and the amount of food by f .

EE static agent

$$o_1 = a_{vr}$$

$$o_2 = a_{mr}$$

$$o_3 = f_{vr}$$

The agents differ in the observations on which they base the decision to mate or not. When the evaluation is static, this means that the observations are evaluated as such, in contrast to relative evaluation where observations are compared with global averages.

2.3 Performance measures

Average population size is used as a measure for the performance of a population. Because not all simulations last 15000 cycles (the maximum), this has to be taken into account too. There are several ways to cope with simulations that do not last 15000 cycles. One of them is to consider a population that has not survived as a population that has zero individuals in the remaining cycles. In this way, the average population size decreases. A disadvantage of this method is that the average population size depends very much on the length of the experiment. Another disadvantage is that an unsuccessful simulation may well have been caused by improper initialisation. "Top" simulations are simulations that last for 15000 cycles. Although statistics have been created for all simulations as well, in this paper only the results of "top" simulations are presented. This choice does not influence the conclusion in Section 4, see www.cs.vu.nl/~wai/projects/dea.

3 EXPERIMENTS

In each DEA experiment, one or more parameters of the system are varied. Within one experiment for each set of parameters fifty simulations have been performed. For example, in the experiment called "time between food injections" the time between food injections was varied from

one to forty-six cycles with steps of five cycles. In that experiment fifty simulations were performed with a time between food injections of one, fifty simulations with six cycles between food injections, and so on.

Many experiments have been done, however, due to space limitations only a representative selection of experiments is presented in this paper. Experiments not included here concern the initial energy of the agents, the maximum age, a maximum number of children, a maximum number of agents, the visual and mating range. The results of those experiments can be found on <http://www.cs.vu.nl/~wai/projects/dea>.

Section 3.1 discusses the experiment "local maximum of agents" which shows that agents overreact on their environment. In section 3.2 the "time between food injections" experiment is discussed. That experiment motivates the main conclusion of this paper, i.e., that PS agents perform worse than EE agents when timing is important. The "randomness of the food distribution" experiment presented in Section 3.3 refutes the hypothesis formed after performing initial DEA experiments that a random distribution of food would be the basis for the difference between EE and PS agents.

3.1 Local maximum of agents

Initial experiments showed that PS agents outperform EE agents. The hypothesis was that EE agents reproduce too much when there is much food in the neighbourhood. This could be caused by the fact that EE agents can see when there is much food. Simulations have been conducted to test whether this is true. In the simulations a local agent maximum was used, inhibiting reproduction of an agent when the number of agents in its direct surroundings is too high. Without a local agent maximum, EE agents may find a situation with a lot of food a good moment to reproduce, unaware of the fact that all other agents in the neighbourhood are about to do the same thing. When this happens, the number of agents in the neighbourhood would rise sharply. This may cause a food shortage that causes agents to die. In these situations a lot of energy is wasted on reproduction that does not result in a population that is just as large as when the energy would have been used for reproduction at a better moment.

The simulation settings to test this idea the maximum number of agents in the neighbourhood and the size of the neighbourhood were varied. The size of the neighbourhood is determined by a variable called the maximum agent range. Varying the visual range enable the prevention of the maximum agent being larger than the visual range. Below are the results of simulations with a visual range of 7 (default) and maximum agent range of 3 and 4.

FIGURE 2 LOCAL MAXIMUM OF AGENTS WITH RANGE 3

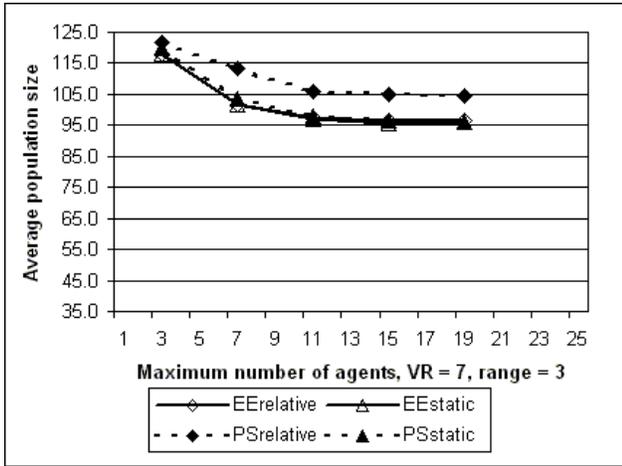
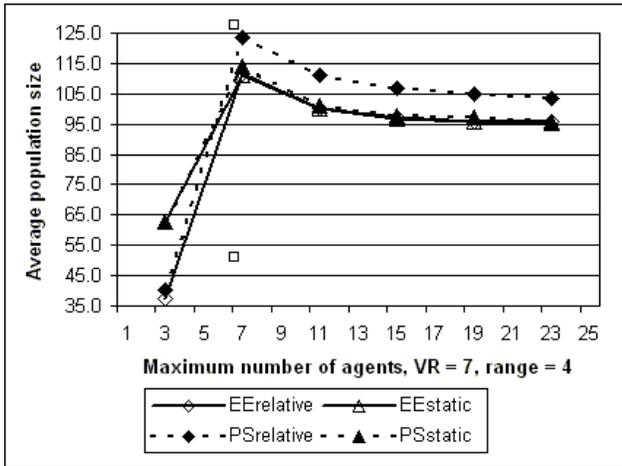


FIGURE 3 LOCAL MAXIMUM OF AGENTS WITH RANGE 4



The experiment shows that preventing reproduction if the number of agents in the neighbourhood is too large has a positive influence on the average population size. This is true, not only for EE agents, but also for PS agents. To measure which type of agents suffers most from this, the best results (highest average population size) of each agent type are compared to result of that type in simulations with default settings. This is the same as measuring which type of agents suffers most from not restricting reproduction. The graphs in Figures 2 and 3 present the best results for each agent type. More simulations have been done, but for clarity only those graphs are shown in which the difference with default simulations is the greatest. These differences are presented in Table 2. For instance, the highest average population size under restricted reproduction for EE relative agents is 118.1, normally this average is 94.6, the difference between the two is 23.5.

TABLE 2 EFFECT OF LOCAL MAXIMUM OF AGENTS

Agent type	Difference with default simulation
EE relative	23.5
EE static	23.7
PS relative	18.8
PS static	20

Table 2 shows that the PS relative agents have the smallest advantage when reproduction is restricted. This means that the PS relative agents suffer least from unrestricted reproduction, whereas EE static agents suffer most.

The graphs show a low average population size when the maximum agent range is 4 and the maximum number of agents is 3. In the graphs with a maximum agent range of 3 this does not occur. This difference might be explained by the fact that a maximum of 3 agents in a range of 4 is far more restricting than a maximum of 3 agents in a range of 3. Restricting the agents more in a setting with a maximum agent range of 3 (by decreasing the maximum number of agents) will probably lead to a worse performance. This idea has not been tested.

3.2 Time between food injections

In this experiment, the time between food injections has been varied from 1 (default) to 46 cycles with steps of 5. The amount of food that is injected is $15 \cdot \langle \text{time between food injections} \rangle$. In this way, the agents receive the same amount of food in all simulations. To find out what the influence of looking into the future is, two additional agents have been created: the EErelativeFuture and EEstaticFuture agent. These agents are the similar to the EE agents, except for their observations:

EE relative future

$$o_1 = a_{vr} - \frac{r_v}{s} \cdot a_{all}$$

$$o_2 = f_{vr} - \frac{r_v}{s} \cdot f_{all}$$

$$o_3 = t_{fi}$$

Where t_{fi} stands for the time until the next food injection.

EE static future

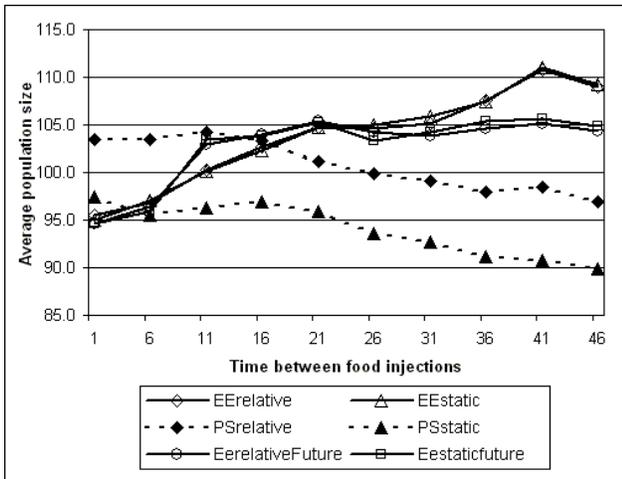
$$o_1 = a_{vr}$$

$$o_2 = a_{mr}$$

$$o_3 = t_{fi}$$

The difference between these agents and normal EE agents is that these agents can observe the time until next food injection. However, they cannot observe the number of agents in mating range, which makes it hard to determine whether differences in results are caused by the fact that future agents can observe the time until the next food injection or by the fact that they cannot observe the number of agents in mating range.

FIGURE 4 VARYING TIME BETWEEN FOOD INJECTIONS



It needed to be investigated whether knowing when the next food injection will come is positive for EE agents. Results of the experiments are shown in Figure 4.

It can be observed that the time between food injections has a positive influence on EE agents, while it has a negative influence on PS agents. Furthermore, the ability to look into the future does not have a positive effect on EE agents. As expected when the time between food injections is small, there is little difference between normal EE agents and their paranormal counterparts. However, when the time between food injections increases, normal EE agents outperform the EE future agents. Therefore, EE agents do not benefit from the ability to look into the future.

From Figure 4 it follows that EE agents outperform PS agents, when timing of reproduction is important. This similar to findings in [2, 7] that a good timing is obtained by observing the environment.

3.3 Randomness of the food distribution

Early experiments were criticised because in the default setting food is injected at random places in the environment. To investigate this criticism, an experiment has been conducted with a world in which food injections vary from totally random to very structured. In the structured setting the food is placed on 9 locations that are equally spaced over the environment, as shown in Figure 5

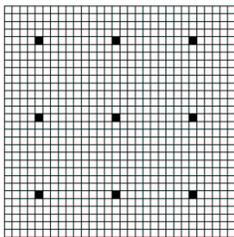
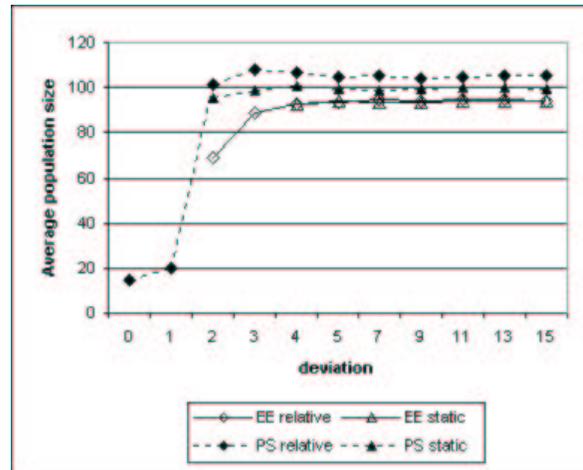


FIGURE 5 ORDERED ENVIRONMENT

The black patches in the world represent the availability food. When the deviation around these points is increased, the food is placed more randomly around the world. Thus, making it possible to investigate the influence of ordering in the environment on the agents' behaviour.

FIGURE 6 FOOD DISTRIBUTION



The more deviation from these points is allowed, the more random the food distribution becomes. Because environment evaluation is outperformed in almost all other experiments it is expected that EE agents would have more difficulty in surviving when isolated from each other.

Figure 6 shows that the performance of all agents decreases when the randomness is decreased (low deviation). This result is more obvious for EE agents than for PS agents.

Compared to the agents of Todd et al. [11], the DEA agents in these experiments react less to diversity with respect to spatial distribution of food. The DEA agents determine the direction of their movements themselves, this is not influenced by evolution, opposed to the agents of Todd, which depend on the environmental structure to adapt their movements.

This experiment shows that randomness of food distribution does not explain the difference between PS and EE agents.

4 DISCUSSION AND CONCLUSIONS

The DEA experiments show that in general partner selection (PS) agents outperform environment evaluation agents (EE). There is one major exception: EE agents outperform PS agents when the time between food injection is increased. This suggests that EE agents are better in timing their reproduction than PS agents, when food is not injected every cycle.

A reason that EE agents perform worse than PS agents when food is injected every cycle could be that EE agents over-react when they see much food. This conclusion is supported by the local agent maximum experiment in which it was shown that when reproduction is inhibited if there are too many agents in the neighbourhood, the EE agents perform better than in simulations without such a restriction. Furthermore the result is greater for EE agents than for PS agents.

From the fact that PS agents perform better than EE agents when food is injected every cycle, it can be concluded that partner selection is also important. An obvious improvement would be to integrate both strategies into one agent; which uses environment evaluation for the timing of reproduction and partner selection for finding a proper mate. However, this was not the main focus of this research.

Translating the DEA results back to nature is notoriously difficult due to the performed abstraction from nature. On the other hand, the emergent behaviour of artificial systems might well help to evaluate theories from biologists. In case of the DEA project the results correspond to the conjecture from research of evolution in nature [2, 7] that environmental clues are used to determine the appropriate time for reproduction.

In both the work of Fletcher et al. [4, 5], and in the DEA project the agents cannot observe the environmental structure completely. The environmental structure consists of a temporal and a spatial structure. The agents in [4, 5] cannot observe the spatial structure. In the DEA experiments, the agents cannot observe the temporal structure.

Comparing the DEA project to Ventrella's work [14, 15] the following two remarks can be made. First, Ventrella's agents select their mates on the basis of physical properties such as colour or body weight. DEA partner selection agents select their mates on the basis of measures that could indicate their fitness, like age. Being old means that the agent is at least capable of surviving.

Secondly, in both Ventrella's work and in the DEA project the agents' preferences with respect to attractiveness are evolved. Both sets of experiments show that agents evolve in such a way that they are not attracted to unfit agents. During the evolution process agents that are attracted to unfit mates produce unfit children and agents with this preference are quickly replaced by agents that are attracted to fit mates.

5 ACKNOWLEDGEMENT

The authors would like to thank Charlotte Hemelrijk for her feedback during the first stages of the DEA project with respect to the possible influence of environment structure on our results and pointers to relevant literature.

6 REFERENCES

- [1] Bedau, M.A., Zwick, M., and Bahm, A., (1995), Variance and Uncertainty Measures of Population Diversity Dynamics, In: *Advances in Systems Science and Applications Special Issue I*, pp. 7—12.
- [2] Carlson, N.R., (2001). *Physiology of Behavior*. Allyn and Bacon Publishing Company (7th Edition).
- [3] Channon, A.D., and Damper, R., (1998), Evolving novel behaviors via natural selection. In: (Adami, C. et al. eds.) *Proceedings of the 6th International Conference on Artificial Life*, MIT Press, pp. 84—88.
- [4] Fletcher, J., Zwick, M., and Bedau, M.A., (1996). Dependence of adaptability on environmental structure. In: *Adaptive Behaviour*, 4(3-4), pp. 275—307.
- [5] Fletcher, J.A., Bedau, M.A., and Zwick, M., (1998). Effect of environmental structure on evolutionary adaptation. In: (Adami, C., et al. eds.) *Proceedings of the 6th International Conference on Artificial Life*, MIT Press, pp.189—198.

- [6] Huang, C.F., (2001), An Analysis of Mate Selection in Genetic Algorithms, In: (Spector, L., et al eds.) *Proceedings of the Genetic and Evolutionary Computation Conference (GECCO-2001)*, Morgan Kaufmann, pp. 766.
- [7] Krapu, G.L., and Reinecke, K.J., 1992. Foraging ecology and nutrition. Ecology and Management of Breeding Waterfowl. University of Minnesota Press, Minneapolis, MN. Northern Prairie Wildlife Research Center Home Page. Version 02FEB99.
- [8] Miller, G.F., and Todd, P.M., (1995). The Role of Mate Choice in Biocomputation: Sexual Selection as a Process of Search, Optimization and Diversification. In: *Lecture Notes in Computer Science*, 899, pp. 169—204.
- [9] Noble, J., (1999). Sexual signalling in an artificial population: when does the handicap principle work? In: *Proceedings of the 5th European Conference on Artificial Life*, Springer-Verlag, pp. 644—653.
- [10] Seth, A.K., (1998), The Evolution of Complexity and the Value of Variability, In: (Adami, C., et al. eds.) *Proceedings of the 6th International Conference on Artificial Life*, MIT Press, pp. 209—220.
- [11] Todd, P., Wilson, S., Somayaji, A., and Yanco, H., (1994). The blind breeding the blind: Adaptive behavior without looking. In: (Husbands, C.D., et al. Eds.) *From Animals to Animats: The 3rd International Conference on Simulation of Adaptive Behavior*, MIT Press, pp. 228—237.
- [12] Todd, P., and Wilson, S., (1993). Environment structure and adaptive behavior from the ground up. In: (Meyer, J.A., and Wilson, S.W. eds.) *From Animals to Animats 2: Proceedings of the second International Conference on Simulation of Adaptive Behavior*, MIT Press/Bradford Books, Cambridge MA, pp. 11—20.
- [13] Todd, P., and Yanco, H., (1996), Environmental Effects on Minimal Behaviors in the Minimat World. In: *Adaptive Behavior*, Vol. 4, Nos. 3 and 4, 1996, pp. 365-413.
- [14] Ventrella, J., (1998). Sexual swimmers: Emergent morphology and locomotion without a fitness function. In: *From animals to Animats*, MIT Press, page 484.
- [15] Ventrella, J., (1998). Attractiveness vs. Efficiency (how mate preference affects locomotion in the evolution of artificial swimming organisms). In: *Artificial Life VI*, MIT Press.
- [16] Wilson, S., (1991). The Animat Path to AI. In: (Meyer, J.A., and Wilson, S.W. eds.) *From Animals to Animats: Proceedings of the first International Conference on Simulation of Adaptive Behavior*, MIT Press/Bradford Books, Cambridge MA, pp. 15—21.