

Evolving Morphological and Behavioral Diversity Without Predefined Behavior Primitives

Peter-Paul Pichler and Lola Cañamero

Adaptive Systems Research Group, University of Hertfordshire
AL10 9AB Hatfield, Herts, UK
p.pichler@herts.ac.uk

Abstract

Virtual ecosystems, where natural selection is used to evolve complex agent behavior, are often preferred to traditional genetic algorithms because the absence of an explicitly defined fitness allows for a less constrained evolutionary process. However, these model ecosystems typically pre-specify a discrete set of possible action primitives the agents can perform. We think that this also constrains the evolutionary process with the modellers preconceptions of what possible solutions could be. Therefore, we propose an ecosystem model to evolve complete agents where all higher-level behavior results strictly from the interplay between extremely simple components and where no ‘behavior primitives’ are defined. On the basis of four distinct survival strategies we show that such primitives are not necessary to evolve behavioral diversity even in a simple and homogeneous environment.

Introduction

The evolution of ‘novel’ behavior by autonomous agents in any simulated system is determined by the predefined components and dynamics of that system. Consequently, the evolutionary possibilities of such a system are necessarily restricted and biased by the preconceptions of the designer. Artificial ecosystems like *Echo* (Holland, 1990), *PolyWorld* (Yaeger, 1994; Yaeger and Sporns, 2006), *LEE* (Menczer and Belew, 1996b,a), or *Geb* (Channon and Damper, 1998) use *natural selection* to overcome one of these biases imposed by the need for an explicit fitness function (artificial selection) in traditional genetic algorithms. All these models vary in the employed level of abstraction and in the details regarding constituents of the agents under evolutionary control (e.g. sensory system, controller, actuation, and morphological properties). However, all of these models resemble each other in that the agents adapt to choose from a predetermined and discrete set of *behavior primitives* which are assumed to be relevant for survival (e.g. *eating, mating, fighting, moving, turning*). This forces the designer of the system to explicitly decide what actions are available and possibly restricts the nature in which they are implemented by the agents.

In our model (Pichler and Cañamero, 2007) actuation is

solely based on movement and reproductive investment. More complex behaviors (e.g. obstacle avoidance, fighting, foraging) are phenomena arising from the interplay between agents and environment. We are interested in *what* strategies arise and *how* they are implemented by low-level interactions of the agent components and its environment in the absence of pre-specified behavior primitives. We think that such an approach further reduces the designer bias and might be more conducive to evolving diverse and adaptive survival strategies.

The results of our simulation show that in such a setting behavioral diversity emerges even in a simple and homogeneous environment. We discuss four different and viable survival strategies and their properties on the level of the individual agent as well as of the whole population.

Virtual Ecosystem

The simulated environment is a space-continuous, time-discrete wrap-around world containing different kinds of objects. All objects in the environment are circular and share certain properties; they have an *energy signature* $e(t)$, a *solidness* ρ and a *radius* r . The energy signature indicates the amount of potentially consumable energy at time t . The solidness determines whether an agent can pass through an object ($\rho = 0$) or whether it collides with it ($\rho > 0$). For agents, radius and solidness are heritable parameters which affect their energy budget in critical ways. Their energy signature is the amount of energy remaining in the world after an agent’s ‘death’ (see next section). Beside agents the environment contains two other types of object: *energy sources* and *obstacles*.

An energy source has a given maximum energy capacity $c > 0$ which defines its initial energy content. If an agent is in contact with an energy source, a certain amount of energy is transferred from the source to the agent and thereby consumed. The energy content of a source cannot fall below zero and ‘grows’ back to its capacity at a constant rate. Energy sources have an energy signature equal to their current energy content, a solidness of zero and a radius equal to their energy signature. Throughout the simulation

they are relocated to random positions with a certain probability. This mechanism was introduced to ‘encourage’ active foraging.

An obstacle is an object with zero energy capacity ($e(t) = 0$) but non-zero solidness. The radius of an obstacle equals its solidness. If an agent collides with an obstacle it is stopped and loses an amount energy depending on its speed and the properties of both objects.

Agent Components

The morphology of an agent is defined by its radius, its solidness and the number and configuration of actuators and sensors along the circumference of its body. Radius and solidness define the mass $m = \rho \cdot r^2\pi$ and the maximum energy capacity $c = \sqrt{m}$ of an agent. The capacity determines the amount and the rate at which the agent can absorb energy from an energy source. It also determines the cost for reproduction (e.g. $0.6 \cdot c$) and influences energy loss (damage) in a collision.

Sensing and Acting

We distinguish two types of sensors; *internal sensors* provide information about the internal variables of the agent and *external sensors* respond to properties of objects in the environment. All sensors function as input nodes to the neural controller network. We define two fixed internal sensors (life energy $l(t)$, reproductive depot $d(t)$) which cannot be removed by evolution. However, they are not necessarily connected to the rest of the network, so it is not predetermined whether or not they are used (see Fig. 2).

External sensors are defined by their position on the body and the type of stimulus they respond to. Each external sensor corresponds to an object property ($e(t), \rho, r$). The information provided by the environment might roughly be thought of as a chemical gradient. The activation a of a sensor s is given by

$$a_s = \sum_{o \in O} \frac{v_o}{d_o^2 + 1} \quad (1)$$

where O is the set of all objects o within a maximum range, v is the value of the respective object property (e.g. solidness) and d is the distance between the sensor and the object.

Every agent has an actuator which regulates reproductive investment. At every time step an energy amount proportional to the activation of that actuator is transferred from the agent’s life energy to its reproductive depot. If this depot reaches a certain threshold, the agent reproduces and an imperfect copy is placed close to it. If an agent ‘dies’ it is replaced by a corpse object with an initial energy content $f(0) = d(t) + 0.1 \cdot c$. Corpses are like energy sources, only their energy decreases (decay) over




Figure 1: Exemplary body of a first generation agent with two sensors (round) and one locomotive actuator (arrow indicates impulse direction).

time. In addition to the reproductive node, an agent can have any number of locomotive actuators. Individually, these work like little jets or flagella, giving an impulse in a specific direction, but combined they can be used to generate more complex movement. A locomotive actuator is defined by its position on the agent’s body and the angle it makes with it (see Fig. 1). This allows us to calculate a rotational and a translational component proportional to the activation of the actuator. The integration over all actuators yields the overall movement of the agent. An actuator is a node in the output layer of the neural controller network.

Neural Controller

Initial agents have few fixed components and no specific functionality. As described above, every agent’s controller network has two internal sensors in the input layer and the reproductive actuator as a node in the output layer. Additionally, initial networks have a small random number of external sensors and locomotive actuators. The two layers are connected by a small random number of links (see Fig. 2). We use nodes with piecewise linear transfer functions and real valued (unbounded) connection weights. The output $N_o(t)$ of a node is given by:

$$N_o(t) = \begin{cases} 0 & \dots & N_a(t) < \theta \\ 1 & \dots & N_a(t) \geq \theta + I \\ \frac{N_a(t) - \theta}{I} & \dots & otherwise \end{cases} \quad (2)$$

where $N_a(t)$ is the accumulated activation of the node, θ is the threshold and I defines a responsive range (slope of the function). The two parameters that define the operating range of a node (θ and I) and the connection weights are randomly initialized and evolved individually for each node and connection respectively. All nodes are arranged in layers and signals travel one layer per time step.

During evolution, both the structure and the parameters of the neural controller networks are freely evolved. Note that in many neuroevolution scenarios (e.g. (Kodjabachian and Meyer, 1998; Stanley and Miikkulainen, 2004)) neural network topologies are evolved to fit specific input and output structure (sensors and actuators). In this model the function and structure of the sensory and actuation systems are completely under evolutionary control. Variability operators

during reproduction may modify all parameters of existing structure and can also add or remove components (sensors, actuators, hidden layers, nodes, and connections) to form arbitrary recurrent networks.




Figure 2: Controller networks have two fixed internal sensor nodes (life energy level $l(t)$, depot level $d(t)$) and the depot node $\Delta d(t)$ in the output layer (left). Additionally, every first generation agent has a (small) random number of sensors, actuators, and connections (right); All parameters are randomly initialized.

Metabolism

The energy budget of an agent is influenced by the properties of its body and its behavior. The base metabolic cost for an agent increases linearly in its mass and in the number of network components. Additional costs are variable and consist of locomotion costs (actuator activation in proportion to mass) and information processing costs (accumulated node activation). These relationships between the agent and the environment defined by the metabolic model shape the dynamics of the system. They create the selection pressures in this artificial ecosystem. All survival-relevant capabilities (sensing, acting, information processing, energy storage) come at an energetic cost. The balance of these aspects should create various trade-offs where agents can follow different strategies to successfully acquire and manage resources and generate a sustained population.

The energy balance of the agents and the resource renewal




Figure 3: Total energy balance of agents and environment.

(energy sources) and decay (corpses) determine the total energy budget of the ecosystem (illustrated in Fig. 3) which is updated every time-step. The ecosystem is not a closed system with respect to energy as energy is added to it and

dissipates via the metabolic consumption of the agents described by the following equations:

$$l_{t+1} = l_t + \Delta e_t - \Delta d_t - C_s - C_{ot} - C_{at} \quad (3)$$

with:

$$C_s = c \cdot \mu_{cc} + \mu_{cc} \quad (4)$$

$$C_{ot} = \sum_{n \in N} a_{nt} \cdot \mu_{co} \quad (5)$$

$$C_{et} = \sum_{e \in E} a_{nt} \cdot \mu_{ce} \quad (6)$$

where l is the life energy level of the agent at time t , Δe is the energy consumed, Δd the energy lost to collision damage, C_s are constant costs (with c the capacity of the agent), C_o are the costs for node activation a over all nodes n in N and C_e are costs for actuator activation a over all actuators e in E (including the investment in reproduction). The μ 's are proportionality constants which were set by trial and error with the goal of balancing the influence of each aspect in a way that each would have a significant and similar impact while still allowing evolution to occur.

The energy content E of an energy source s at time t is:

$$E_s(t+1) = E_s(t) - \sum_{a \in A'} \Delta e_a^s(t) + \mu_g \quad (7)$$

where a is an agent in the set A' of all agents which have consumed energy from source s at time t . The energy content of a source cannot be negative. This equation also holds for corpses if the constant growth rate $\mu_g > 0$ is replaced by a decay rate $\mu_d < 0$.

Reproduction

There are many possible ways to define a reproduction criterion in a foraging scenario like the one presented here. Two straightforward ideas are either a life time dependent criterion or using the life energy of an agent (see e.g. (Bedau et al., 1992)). Here, agents would periodically reproduce after a certain number of time steps or whenever their energy level reaches a specified threshold. However, solely 'optimizing' individual longevity disables survival strategies with short individual life times and thereby excludes potentially interesting dynamics like persistence vs. progeny trade-offs (Polani et al., 2006). The same is true when using the energy threshold as the single criterion; this strips the agent of much of its autonomy on how to manage the acquired resources. Using the reproductive actuator we have a reproduction criterion which gives the agents full control over *when* and to *what extent* they invest in reproduction. Whenever this node is activated an amount of energy proportional to the activation is transferred from the agent's life energy to its reproductive depot. Once this depot reaches a certain threshold, an imperfect copy is

placed close to the agent. Reproduction in our model is strictly asexual. Mutation operators exist to modify all body properties and the topology as well as all parameters of the controller network. While there is no final consensus about what is the best way to encode neural networks for artificial evolution it has been shown repeatedly (see e.g. (Stanley and Miikkulainen, 2002) or (Seys and Beer, 2006)) that the encoding has a crucial impact on the evolvability of the system. Keeping this in mind we presently use no ‘genetic’ encoding and all mutation operators are performed directly on the agent’s object structure (this is equivalent to a direct encoding scheme).

Adaptation and development in this experiment occurs solely on an evolutionary scale through reproduction. Agents do not change or adapt during their lifetime. However, change on an evolutionary scale can only happen if a *turnover of generations* exists. In a classical genetic algorithm this turnover is an inherent property which is explicitly enforced by the design of the algorithm itself. In our model (and other models based on natural selection) this turnover of generations is to some extent an emergent property of the dynamics of the system. Because reproduction is ‘optional’ it is in some sense an adaptation itself. Agents have to actively invest their *life energy* into creating offspring and doing so jeopardizes their own survival because the invested energy is no longer available to them and reproducing creates a direct competitor in the vicinity. A first intuition might suggest that this would eventually lead to zero investment in reproduction. In this case evolution would cease to happen or, in fact, never happen at all. On second thought, however, it is clear that in a dynamic based on natural selection the notion of selecting *for* zero reproduction is contradictory as reproduction is the very vehicle of selection. Additionally, in an environment where individual survival is to some degree dependent on chance and thus effective immortality is unachievable, an infertile population is unsustainable and inevitably doomed. Randomly created agents are more often than not unable to survive for any length of time, let alone spare enough energy to reproduce if they even do so at all. To guarantee a certain number of agents in the environment we use a mechanism similar to (Yaeger, 1994). The *minimum enforced agents mechanism* (MEAM) creates new random agents whenever the total population size falls below a given threshold. Therefore, it guarantees that there are always agents present in the environment but becomes inactive once agents reproduce and successfully establish a sustained population of a certain size. Population size is therefore not fixed or constant, but depends on the environment and the properties of the evolved agents (see next section). To track the existence of a generational turnover we assign a *phylogenetic generation* (PG) to each agent. Agents created by the MEAM have a PG of zero, their offspring a PG of one, and so on. Evolution only occurs if this number

increases.

Experiment and Results

To obtain the results discussed in this paper the simulation was run in relatively small 100x100 unit arenas (minimum agent size is 0.1 units) with 35 energy sources and 35 obstacles. Energy sources had an energy capacity of 1.0 and obstacles a solidness of 1.0. Objects were randomly placed in the environment following a uniform distribution. We repeated the simulation 85 times using different random seeds for the random number generator which determines object placement, initial agent configuration and all mutation operators. The minimum enforced number of agents was 15 in all 85 runs. Since in this setup there is no obvious ‘convergence point’, simulations were run until the average PG of a population was above 500 or a set maximum time was reached (80 hours). From each of the 76 ‘successful’ runs (where a sustained population was established) a sample of the first 100 agents of $PG \geq 500$ was taken.

Behavior and Morphology

In 76 out of 85 runs the MEAM eventually established a sustained population and evolution could occur. Actual computation time to reach this point depended greatly on a number of factors: the moment a sustained population was established, the average population size, the complexity of the average controller network, and the average lifetime of the individual agents. While in some runs a sustained population was established almost immediately, in 9 runs it did not happen at all before the maximum time was reached. These runs were discarded. A general observation was that all populations were quite homogeneous within a single run. On reason for that is that all agents within the population of one run were ultimately descendants of one respective *founder* agent which spawned the initial population. Other possible reasons are that the environments were rather small (an agent could travel ‘around the world’ frequently during its lifetime) and both obstacles and energy sources were uniformly distributed. In the following sections we will describe some of the evolved agents, their behavior and their morphologies (for illustrative examples that convey the nature of the evolved strategies much better than words we kindly refer the reader to the videos on the first author’s website¹). All agents in this experiment exhibited base movement (movement in the absence of stimulus). For the first part of the analysis of the results we distinguish three basic evolved behavior patterns solely by observable behavior:

- *Energy response*: agents show *some* response (e.g. slowing down) in the presence of or on contact with an energy source.

¹<http://homepages.feis.herts.ac.uk/~pp6bs/>

	ER	EA	OA
Drifter	yes	no	no
Forager	yes	yes	no
Avoider	yes	no	yes
Allrounder	yes	yes	yes

Table 1: Classification of agents by the three observable behavior patterns: ER (energy response), EA (energy approach), and OA (obstacle avoid).

- *Energy approach*: agents change direction and actively try to approach an energy source.
- *Obstacle avoidance*: agents change their behavior in the presence of an object of non-zero solidness.

The definition of these behavioral patterns is intentionally careful. If an agent changes its behavior in response to an obstacle, it might do it in a way that will generally increase the probability of avoiding a collision. However, those mechanisms are not perfect and in some situations the behavioral change of the agent might actually cause it to hit the obstacle even harder than without any change. Because behavior is the result of the interaction between the agent and the environment (Beer, 1995), no observer would speak of the resulting behavior as obstacle avoidance if the agent actually makes the impact worse.

We have classified the agent strategies in four basic kinds, based on the three behavioral patterns identified above (see Tab. 1). Overall, agents of the same class share essential behavioral tendencies, even though they vary in the details of their implementation. Figure 4 shows the morphological properties of the four agent classes, categorized by behavior patterns, and Fig. 5 shows differences between agent categories on the population level. It is interesting (though not surprising) to note that even though the categorization was done solely on the basis of behavioral observations it is nearly perfectly reflected in the body properties of the agents. As could be expected it turns out that if both body and controllers are evolved as a functional unit one cannot discuss one without the other. The evolutionary dynamics shape the complete agent and adapt it to a certain survival strategy.

Drifters exhibit relatively fast base movement using their (usually) single functional locomotive actuator. With a single actuator an agent cannot change its direction, it can only modulate its speed. Consequently, drifters can neither avoid obstacles nor can they actively approach an energy source. Instead, they modify their speed in the presence of an energy source. This is achieved either by ‘monitoring’ their life energy supply and stop moving if it exceeds a certain threshold, or by using energy sensors to measure the energy concentration of the environment. Whenever the energy concentration

is high they slow down or come to a complete stop. Drifters are typically very small and light-weight (see Fig. 4). Their life span is comparatively short but their population size is larger than that of all other types (see Fig. 5). Drifters usually only have one sensor, one (functional) actuator and minimal networks to control their extremely simple behavior. In many simulation runs, the first sustained population consists of drifter-like agents. Sometimes they evolve into other types, but often a relatively stable drifter population establishes itself where only the morphological properties are further refined to suit this strategy. It is worth noting that even this simple strategy requires a fair amount of adaptation to first acquire and then ‘calibrate’ the required sensory and actuation system. No viable strategy emerged where output was constant (e.g. comparable to ‘always go forward and kill’ reported in (Channon and Damber, 1998)).

Foragers have base movement, change their behavior in the presence of an energy source, but do not avoid obstacles. In their simplest form, a single energy sensor and one actuator placed roughly opposite the one responsible for base movement are sufficient to perform successful approach behavior. The translational component of the base actuator is counteracted by a usually slightly tilted second actuator. This results in an inward spiraling movement dependent on the strength of the sensory stimulus. However, the experiments show that usually two energy sensors and a larger number of actuators are used to implement this behavior. Also the actual behavior resulting from the agents’ actions and its robustness vary from population to population and over evolutionary time. Some agents will always manage to approach an energy source within their sensor range while others may only succeed if they are approaching from a particular side. Another difference is how well an agent is able to keep contact once it has approached the energy source. While some agents spend most of the time ineffectively circling around an energy source, others can perfectly center themselves over them and remain there until the source is either fully consumed or disappears.

While drifters usually minimize their body size to the lower bound of 0.1, foragers almost consistently have a size of about 0.3 (see Fig. 4). Some foragers also increase their solidness instead of their size. Both adaptations lead to higher capacity but also higher movement costs. Foragers have to find energy more often than drifters but also consume energy sources more efficiently.

Avoiders follow a somewhat surprising strategy. They are the only agents that completely abandon energy perception through external sensors. Avoiders exhibit base movement and obstacle avoidance. The different populations responded to contact with energy sources in different ways. In all cases the resulting behavior can be explained by the internal sensor for the agent’s life energy level. In the first case

Morphological Properties

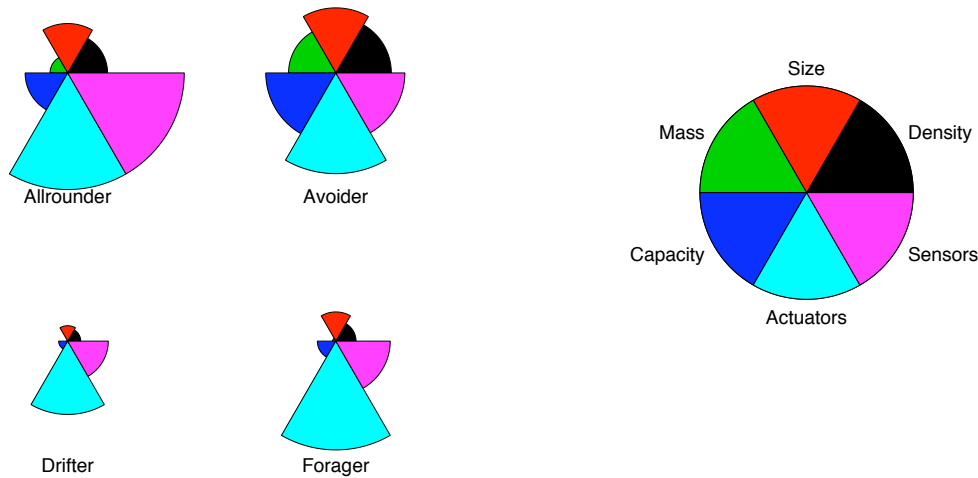


Figure 4: Morphological properties (all normalized to 1) of evolved agents (PG 500) categorized by behavior patterns.

the base actuator of the agent is inhibited once its life energy level exceeds a certain threshold and the agent stops on top of the energy source. In the second case the same trigger activates the actuator used for obstacle avoidance causing the agent to start moving on a perfectly circular trajectory. In both cases the agent (at least partly) consumes the energy source without directly sensing its presence. The respective behavior patterns persist even if the energy source disappears until the life energy level drops below the triggering threshold. Avoiders have slower base movement than other agents. This seems to be an adaptation to their increased weight and their consumption strategy as there is a considerable delay between first contact with the energy source and the life energy reaching the needed threshold to trigger the agent's response. The observed avoiders are bigger than foragers and have a higher solidness. The increased solidness gives them a much larger capacity at a medium risk because of their obstacle avoidance capabilities.

Allrounders are agents which exhibit all three behavior patterns. Basically they are the same as foragers with the added ability to avoid obstacles. Their foraging behavior is the same and they can sometimes evolve from forager agents. However, they tend to have a higher capacity than basic foragers. Most of the evolved allrounders achieve this by increasing the solidness value. As with avoiders the risk of increasing the solidness is lowered by the ability to avoid obstacles. Allrounders (as can be expected) have the most complex networks and the most sensors and actuators. They

also have the smallest population sizes.

General Properties To show that behavioral diversity emerges even in simple and uniform environments we have only presented four survival strategies. However, it is worth noting that changing only the concentration of obstacles and energy sources can lead to completely different behavioral strategies. We will mention one observed type of agent because of their radically different approach. This strategy appeared in environments with high concentration of both obstacles and food sources. A high concentration of obstacles 'penalizes' movement early in evolution when agents are not yet well adapted (by either being light-weight or by avoiding the obstacles). There, agents can be nearly or completely sessile. These agents exhibit no base movement at all. They remain stationary until an energy source appears within their sensor range. Once in range, they quickly approach the energy source, center themselves over it and remain there. These agents have much larger bodies and simpler controller networks than mobile agents. Larger size consumes a lot of energy when moving but it also increases the maximum energy capacity of the agent. A larger agent which does not move can survive longer without consuming energy.

More generally, however, selection seems to favour small and light-weight agents that exhibit some base movement early in the evolution. This is further optimized if agents follow the drifter strategy. Agents with an active foraging strategy (foragers and allrounders) are slightly larger and agents

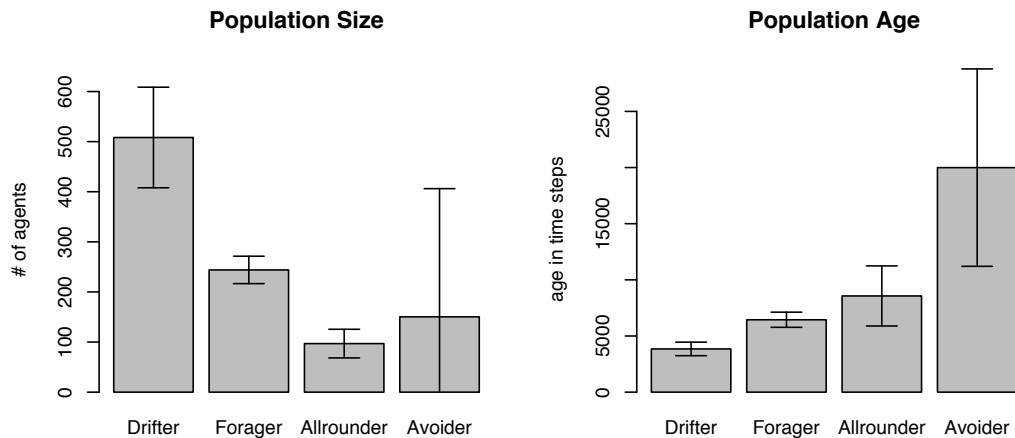


Figure 5: Average population size (left) and average population age (right) of evolved agent populations categorized by behavioral strategies. Error bars show the 95% confidence interval. The large bars for avoider populations are due to the small sample size (3).

with slow base movement are even larger still to increase their energy capacity. All agents without collision avoidance minimize solidness. Agents with collision avoidance often increase solidness and size to increase their energy capacity. Sensors are effectively restricted to the required minimum while actuators seem to accumulate even if they are not used efficiently or not at all (see Fig. 4).

Reproductive strategies are very hard to analyze in detail as they can only be understood by analyzing the dynamics of each agent's network. Supporting the rationale behind our reproductive criterion which gives the agents control over *when* they invest in offspring, constant reproductive activity (irrespective of internal and external circumstances) did not emerge as a viable strategy in a single sustained population. However, most agents follow simple reproductive strategies or combinations thereof; these can roughly be summarised as follows: Invest in reproduction if energy is present, otherwise don't. There are different ways to achieve this. The most commonly used is a positive correlation between the energy sensors and the actuator for the reproductive depot. Alternatively, the activation of the reproductive depot is positively correlated to either the internal energy level or the activation of a locomotive actuator used for foraging. Many agents use a combination of these strategies. Additionally, often a negative correlation between a solidness sensor (or an actuator used for collision avoidance) and the reproductive activity exists.

Discussion & Conclusion

We have shown data of four distinct behavioral strategies evolved in a virtual ecosystem. The different types of agents

evolved 'high level' behaviors (foraging, obstacle avoidance) without a discrete set of predefined behavior primitives and without other pre-defined functionality or structure. All behavior is the result of the agents interacting with the environment via a very simple but versatile locomotion model. The evolution was done in an artificial ecosystem by natural selection and both neurocontrollers as well as morphology (size, solidness, sensory and actuation structure) of the agents were freely evolved. Based on the results of these more general experiments we are satisfied that this approach is very capable of evolving diverse behavior while further reducing the need preconceive necessary action possibilities the agents might need to survive under different environmental conditions. To keep evolved strategies comparable we have only used a small and homogeneous environment in this experiment. A possible extension of the presented experiment is to investigate the impact of more variable environments on the evolution of survival strategies.

While we think that replacing discrete behavior primitives by our simpler actuation model in combination with the proposed reproduction criterion is more conducive to the evolution of diverse behavior, it is also clear that such a reduction of the set of predefined biases is not possible or even desirable *ad infinitum*. Apart from obvious computational complexity considerations the actual goal of the simulation has to be considered. We tried to create an evolutionary setting which is flexible enough to allow the evolution of distinctly diverse and non-trivial agent strategies. In other situations a different set of biases might be appropriate. One main message of this paper is that, also when using natural selection in an ecosystem scenario, one has to be aware what biases

are built into the system and how they affect the simulation. One avenue of future research will consist of a comparative study about how different reproductive criteria influence the evolved diversity of agent strategies.

Another future aim of this project is to investigate the potential emergence of phenomena comparable to basic affect in natural organisms. Basic affect in this context includes individualistic affect like approach-avoidance, arousal and agonistic affect, as well as prosocial affect like cooperation. These phenomena are thought to be the physiological bases for higher level affect (as e.g. described in (Buck, 1999)).

We are currently extending our ecosystem model to include the possibility to evolve simple neuromodulatory mechanisms which are used in animals to support affect. Similar to neuromodulation these mechanisms would allow the neurocontrollers to regulate whole groups of neurons as opposed to the direct synaptic transmission in standard neural networks. Therefore, in the next set of experiments we will investigate if providing this possibility will lead to the evolution of agents that exhibit properties normally ascribed to such basic affect. Targeted results of these experiments include changes in foraging behaviour depending on the life energy level (arousal) or flexible weighting in approach-avoidance conflicts (e.g. approaching energy source close to an obstacle only in certain situations). If such mechanisms are successfully evolved we expect agents to develop more flexible behavior strategies which are also more robust to changes in the environment. We also hope to be able to draw some conclusions about the necessary conditions and origins of functionally similar processes in real organisms.

Acknowledgment

Peter-Paul Pichler is funded by a research scholarship of the University of Hertfordshire.

References

- Bedau, M. A., Norman, A., and Packard, H. (1992). Measurement of evolutionary activity, teleology, and life. In Langton, C., Taylor, C., Farmer, D., and Rasmussen, S., editors, *Artificial Life II*, pages 431–461, Redwood City, CA. Addison-Wesley.
- Beer, R. D. (1995). A dynamical systems perspective on agent-environment interaction. *Artif. Intell.*, 72(1-2):173–215.
- Buck, R. (1999). The biological affects: A typology. *Psychological Review*, 106(2):301–336.
- Channon, A. D. and Damper, R. I. (1998). Evolving novel behaviors via natural selection. In Adami, C., Belew, R. K., Kitano, H., and Taylor, C., editors, *Proceedings of the 6th International Conference on Artificial Life*, pages 27–29, Cambridge, MA, USA. MIT Press.
- Holland, J. H. (1990). Echo: Explorations of evolution in a miniature world. In Langton, C. G., Farmer, J. D., and Doyne, J., editors, *Second Conference on Artificial Life*. Addison-Wesley.
- Kodjabachian, J. and Meyer, J. (1998). Evolution and development of neural networks controlling locomotion, gradient-following, and obstacleavoidance in artificial insects. *IEEE Transactions on Neural Networks*, 9(5):796–812.
- Menczer, F. and Belew, R. K. (1996a). From complex environments to complex behaviors. In *Adaptive Behavior 4*, pages 317–363.
- Menczer, F. and Belew, R. K. (1996b). Latent energy environments. In Belew, R. K. and Mitchell, M., editors, *Adaptive Individuals in Evolving Populations: Models and Algorithms*, pages 191–208, Reading, MA. Addison Wesley.
- Pichler, P.-P. and Cañamero, L. (2007). An evolving ecosystems approach to generating complex agent behaviour. In *Proceedings of the First IEEE Symposium on Artificial Life*.
- Polani, D., Nehaniv, C. L., Martinetz, T., and Kim, J. T. (2006). Relevant information in optimized persistence vs. progeny strategies. In M.Rocha, L., Bedau, M., Floreano, D., Goldstone, R., Vespignani, A., and Yaeger, L., editors, *Proc. Artificial Life X*.
- Seys, C. W. and Beer, R. D. (2006). Effect of encoding on the evolvability of an embodied neural network. In *GECCO 2006 Workshop Proceedings, Workshop on Complexity through Development and Self-Organizing Representations (CODESOAR)*.
- Stanley, K. O. and Miikkulainen, R. (2002). Efficient evolution of neural network topologies. In *Proceedings of the 2002 Congress on Evolutionary Computation (CEC '02)*.
- Stanley, K. O. and Miikkulainen, R. (2004). Competitive coevolution through evolutionary complexification. *J. Artif. Intell. Res. (JAIR)*, 21:63–100.
- Yaeger, L. (1994). Computational genetics, physiology, metabolism, neural systems, learning, vision and behaviour or polyworld: Life in a new context. In Langton, C. G., editor, *Artificial Life III*, volume XVII, pages 263–298. Addison-Wesley.
- Yaeger, L. S. and Sporns, O. (2006). Evolution of neural structure and complexity in a computational ecology. In Rocha, L. e. a., editor, *Artificial Life X*, Cambridge, MA. MIT Press.