An Evolving Ecosystems Approach to Generating Complex Agent Behaviour

Peter-Paul Pichler
Adaptive Systems Research Group
School of Computer Science
University of Hertfordshire
Hatfield, Herts, AL10 9AB, UK
Email: p.pichler@herts.ac.uk

Lola Cañamero
Adaptive Systems Research Group
School of Computer Science
University of Hertfordshire
Hatfield, Herts, AL10 9AB, UK
Email: l.canamero@herts.ac.uk

Abstract— We propose an evolving ecosystem approach to evolving complex agent behaviour based on the principle of natural selection. The agents start with very limited functional design and morphology and neural controllers are concurrently evolved as functional wholes. The agents are 'grounded' in an increasingly complex environment by a complex model metabolism and interaction dynamics. Furthermore, we introduce a novel criterion for evaluating differential reproductive success aimed at maximising evolutionary freedom. We also present first experimental results suggesting that this approach may be conducive to widening the scope of artificial evolution for the generation of agents exhibiting non-trivial behaviours in a complex ecosystem.

I. Introduction

Attempts to apply methods of evolutionary computation to create artificial agents with non-trivial behavioural complexity often run into several problems. One major problem when using classical genetic algorithms arises from the fact that they require a fitness function to guide them to the desired results. However, as has frequently been pointed out [3], in many contexts it is notoriously difficult to design such a function. Especially if a solution is hard to evolve, it is crucial that the fitness function contains not only measures that assess the desired evolutionary outcome, but to permit successful evolution, it must also provide a meaningful gradient along the way. This usually requires a lot of experience and tweaking on the part of the experimenter and may in some cases well be as complex as solving the task itself. Several ideas [5], [4] have been brought forward, to address this bootstrapping problem of evolving complex agent behaviour. An obvious way is to relocate the starting point of evolution by inserting certain knowledge into the system. For instance, this can be done by 'hard-wiring' some basic behaviours [1], or otherwise structurally or functionally predisposing the agents in some advantageous way.

Another proposed method of sidestepping the problem is to aid the evolutionary process by incrementally raising the difficulty by partitioning the evolutionary task into manageable sub-tasks. Once these sub-tasks have been isolated, evolution is either iterated with a hierarchy of fitness functions or the complexity of the environment is incrementally increased. This method called incremental evolution was shown by [5] to perform well if a successful partition can be found. However,

[6] suggest that this may not generally be the case.

To tackle the bootstrapping problem, the aforementioned methods share the policy that if a problem is 'too big' for evolutionary computational methods to solve, it is made smaller by adding knowledge to the system. While this may certainly lead to the evolution of more complex agent behaviour and thereby better solve a given problem, it does not necessarily render the evolutionary process itself more powerful.

In artificial life, the focus of study can be shifted from the desire to reach a predefined goal to the question of how natural evolution actually avoids these problems. One step in this direction is to get rid of the fitness function. Several models have been proposed (e.g. see [1], [2], [4], [7]) that stress the need to model organisms and environment as one ecosystem and move from explicit fitness functions to implicit ways to assess reproductive success. However, for varying reasons, most of the conducted experiments have remained intentionally simplistic with respect to overall system dynamics.

We believe that an additional way to smooth the evolutionary gradient is to concurrently evolve all constituents of the agents as a functional whole. The asymmetry in complexity that may result from introducing knowledge in only part of the system may well be another cause for difficulties. Therefore, what may help the evolutionary process if the goal as well as the route to the goal are known, may hinder evolution if this is not the case. Our aim is to test the feasibility of an evolving ecosystem approach if very little 'innate intelligence' [1] is pre-specified and the system dynamics allow for highly non-trivial adaptations. We here propose a model for an artificial evolution of complete virtual organisms which may be more conducive for generating non-trivial results.

II. A Model for an Evolving Ecosystem

Our model implements a classic foraging scenario, where an agent's success is, to some extent, dependent on its ability to acquire and efficiently manage a certain resource which we call *energy*. Energy is an exhaustible but renewable resource and all system properties and their dynamics are in some way connected to it. The flux of energy through the system is the driving force of evolution. The dependencies between energy and the elements of the system have been modelled to allow

for complex interactions and diverse adaptations that can hopefully allow the generation of diverse survival strategies.

Our model of an evolving ecosystem works akin to the principle of natural selection. Selection pressures arise from the interaction of the agents with the environment and the direct competition between the agents themselves. Once a sustained population occurs, the level of free energy in the environment decreases drastically while competition for these resources increases. This gradually increases selection pressure for more efficient survival strategies and environmental complexity without requiring the experimenter to define the right time or nature of these transitions. We think that this gradual increase, combined with the concurrent evolution of all agent constituents can to some extent alleviate the bootstrapping problem and widen the scope of artificial evolution.

However, our model does not fulfill the criteria for an open-ended evolutionary process in the sense of [8]. Especially, our agents are separate from the environment by external definition and they are not self-replicating. Therefore the need for a criterion which regulates differential reproductive success remains. We think that to optimally exploit the creative potential of evolution the choice of this criterion is crucial. Otherwise it is still possible to end up with dynamics that simply optimise one aspect.

We have tried to devise a simple fitness criterion which will introduce little unintended bias while still providing a gradient for successful evolution. Every agent possesses a dedicated node (*depot node*) in the output layer of its neural controller network. At each time step, an amount of energy proportional to the activation of the node is transferred from the agent's 'life' energy to a *reproduction depot*. If this reproduction depot reaches a certain threshold (proportional to its capacity) the agent reproduces. Note that while this might appear similar to [1] there are fundamental differences. Using this dedicated depot which has to be actively filled with surplus energy the agent gains full control over *when* and *how much* energy it invests in offspring. Therefore, not only is there a direct cost (transferring the energy) involved in creating offspring, but also, doing so will create a direct competitor for available resources. Hence, reproduction in our model is 'optional' and the selective pressure does not arise from a global fitness function but from local and direct competition in a shared environment.

Reproduction in our model is asexual. If an agent reproduces, an imperfect copy is placed near the parent. Mutation operators exist for modifying each of an agent's parameters along with operators for deleting existing, or adding new parts (sensors, actuators, neurons, synapses). We do not use any 'genetic encoding' and all mutation operators work directly on the agent's object structure.

The major challenge in creating such a model lies in modelling the generic agent components and their interaction dynamics with the environment. Once that is done, the 'evolutionary algorithm' itself or, respectively, the main

```
do forever
    all objects ← update()
    for all agents do randomly
        if energy > 0
            agent ← update()
            if depot > R_T
                agent ← reproduce()
        else add corpse ← agent
    while #agents < A_min
        createAgent()
```

Fig. 1. Pseudocode showing the main program loop of the evolutionary system. At every timestep all non-agent objects are updated. Next, agents are either updated and reproduce if their reproductive *depot* exceeds their reproductive threshold R_T, or replaced with a corpse object if their energy is less or equal to zero. Finally, if the total number of agents is lower than a set minimum A_{min}, an appropriate number of random agents is created.

program loop (see fig. 1) is extremely simple; solely updating the environment and agents. The only mechanism that interferes with the system is what we call the *minimum enforced agents mechanism* (MEAM). The MEAM ensures that a specified minimal number of agents is always present in the environment. Once the size of the agent population grows larger through reproduction and remains so, this mechanism has no more effect unless the population size falls below that threshold again.

III. ENVIRONMENT

The experimental environment is a continuous two-dimensional toroidal world that contains a number of different objects. All objects are circular and carry information in the form of a currently three dimensional property vector. These three *information channels* represent *energy content*, *solidness*, and *size* of an object. Even though these three channels are provided by every object, the way in which those values are calculated differs slightly between the different types of objects. In total, we distinguish between four types:

Energy sources are objects without solidness but non-zero energy content. They are initialised to a random energy capacity $c \in (0, C_{max}]$, where C_{max} is a maximum capacity globally set by the experimenter. If consumed by an agent, an energy source re-grows at a constant rate. The energy content of an energy source corresponds to its energy level at time t. The size of an energy source equals its energy content. Because it has no solidness, agents can pass through an energy source unhindered. In the present setup energy sources are relocated to a new random position with a certain probability.

Obstacles are objects without energy content but a non-zero solidness. The size of an obstacle equals its solidness value. Each obstacle is initialised with a random solidness value $\in (0, 1]$. If an agent collides with an obstacle the agent is stopped and receives a certain damage (loss of energy, see section V-B). In the current setup obstacles are static and do not change their position throughout an evolutionary run.

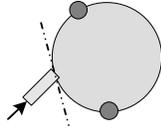**Agents** are just other elements of the environment that can

Fig. 2. Exemplary body of a first generation agent with two sensors (round) and one actuator (arrow indicates impulse direction).

be sensed through the same channels as can every other object. However, unlike with other object types, both size and solidness are genetically specified non-zero parameters. Additionally, the energy content channel in the case of an agent does not refer to its current energy level, but to the value of its reproductive depot (see section IV-A.2).

Corpses are 'dead' agents. A corpse is a food source with the energy capacity of the agent's reproductive depot at its time of 'death'. However, its energy content decreases over time and it is deleted when the energy reaches zero to prevent cluttering of the environment over the course of the simulation.

IV. AGENTS

The agents evolved in our system consist of a body and a controller. This traditional distinction is upheld mainly for explanatory purposes, since in our simulation both parts overlap and concurrently evolve as a functional whole. The initial agents have only very little 'innate intelligence' [1] or functionality. Instead, all aspects of the agents were designed to provide comparatively high levels of evolutionary freedom while at the same time reducing bias introduced by predefined structure or functionality.

A. Morphology

The morphology of an agent is defined by the properties of its virtual body and the number and position of its sensors and actuators. We use a simple circular shape for the agent's body but both the arrangement of sensors and actuators as well as the configuration of other body properties are completely open to evolution. Thus, at least theoretically, a vast number of different morphologies is possible.

1) Body Properties: A circular shape was chosen for the basic body plan because it is generally considered the most neutral form with respect to symmetry; with a circular body the agents have no predefined direction and are rotationally symmetric. Preferred direction of movement and possible symmetry properties are evolutionary emergent and determined by the placement of *external sensors* and *actuators*. Both external sensors and actuators can be situated anywhere along the circumference of the agent's body.

To avoid confusion we emphatically call attention to the fact that it was not our aim to implement our model in a realistic or transferable fashion. Wherever it was easily possible, the dynamics resemble our real world intuitions, but in principle our system implements an arbitrary and independent virtual world. However, each of these properties have different metabolic consequences for an agent, discussed in section IV-C.

The free parameters are initially set to random values within their defined range.

- *Size:* The size $s \in [0, \infty)$ defines the radius of the body.
- *Solidness:* The solidness $\rho \in (0, 1]$ is a free parameter.
- *Mass:* The 'mass' m of an agent is a function of its size and solidness with $m = \rho \cdot \pi s^2$
- *Energy Capacity:* The energy capacity $C = \sqrt{m}$ is the maximum amount of energy an agent can accommodate.
- *Maximum Absorption:* The maximum amount of energy $\Delta\epsilon(t)$ an agent can absorb from an energy source per timestep is given by:

$$\Delta\epsilon(t) = \min(c_\epsilon(1 + \pi s^2), C - E(t)) \qquad (1)$$

where c_ϵ is a proportionality constant and $E(t)$ is the energy level of the agent at time t.

2) Sensors: We provide two predefined groups of sensors which can help the agents make use of different sources of information. Generally, sensors act as input nodes to the neural controller network of the agent. The first kind are *internal sensors* which provide information about the agent's own internal variables. There is one internal sensor for the current energy level and one for the current level of the reproductive depot. Internal sensors are fixed parts of the input layer of the initial controller networks. Even though they cannot be removed by evolution, they are not necessarily connected to the rest of the network, thereby it is not predetermined whether or how they are used.

The second basic kind of sensors are *external sensors*. As the name suggests these sensors allow an agent to perceive information present in the environment. The environment provides sensory information on three different channels (energy content, solidness, and size). In order to make use of these signals an agent must evolve the appropriate sensors. An external sensor of type i responds to the i-th channel of an object. The information provided by a single channel is one dimensional and can roughly be though of as a chemical gradient. The response r of the i-th channel of an object o depends on its value $v_{i,o}$ and the distance d_o between the sensor and the object. It is given by

$$r_i = \frac{v_{i,o}}{1 + c \cdot d_o^2} \qquad (2)$$

Hence, the total activation of a sensor S of type i follows from the accumulation of the responses of all objects within a maximum range.

$$S_i = \sum_{o \in O} \frac{v_{i,o}}{1 + c \cdot d_o^2} \qquad (3)$$

where O is the set of all objects within said maximum range and c is a proportionality constant.

The relationship shown in (3) leads to a sensory response landscape which is further illustrated for a single channel in fig. 3. Additionally, as has been noted in section III, in the case of another agent, the energy channel corresponds to the agent's energy depot instead of its current energy level. The intuition behind this choice is that the energy channel
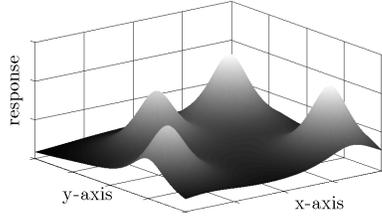
Fig. 3. Sensory response landscape for a given channel in an environment with four objects. The peaks represent the centres of objects where the response is highest.

Fig. 4. Controller networks have two fixed internal sensor nodes (energy level e, depot level f) and the depot node Δf in the output layer (left). Additionally, every first generation agent is initialised with a (small) random number of sensors, actuators, and connections (right); All parameters are randomly initialised.

contains information about '*potentially consumable energy*'. The energy of the corpse of an agent relates to its energy depot level at the time of death.

3) Actuators: Actuators are a predefined structure to enable the evolution of locomotion in the agents. They are very simple in order to allow basic forms of locomotion but can be arbitrarily combined to implement more complex movement strategies. As with sensors, actuators are nodes in the neural controller network, however, they differ by being located in the output layer. Additionally, actuators are defined by their position on the agent's body and the angle formed by the tangent of the body at that point (see fig. 2). Hence, the influence of an actuator is fully described by its position, its direction, and the strength of its activation. This results in a movement vector $\vec{v}$ for each actuator:

$$\vec{v} = \begin{pmatrix} -\sin(\Phi + \alpha + \beta) \\ \cos(\Phi + \alpha + \beta) \end{pmatrix} \qquad (4)$$

where Φ is the rotation of the agent, α is the position of the actuator on the agent's body ($0 \leq \alpha < 2\pi$), and β is the angle of the actuator ($0 \leq \beta < \pi$) with respect to the tangent of the agent's body at its position. This movement vector $\vec{v}$ is normalised and scaled with the activation of the actuator to yield the final movement.

This allows us to calculate a translation and rotation component for every actuator corresponding to its activation. Thus, the overall movement of an agent can be obtained by integrating over all actuators. The effects of multiple actuators on the position $\vec{P}$ and the rotation Φ of the agent at time step t are as follows:

$$\vec{P}(t+1) = \vec{P}(t) + \sum_{i}^{|A|} \frac{\vec{v_i}}{|\vec{v_i}|} \cdot c \cdot o_i \qquad (5)$$

$$\Phi(t+1) = \Phi(t) + \sum_{i}^{|A|} c \cdot o_i \cdot \cos(\beta_i) \qquad (6)$$

where A is the set of actuators, o_i is the activation of the i-th actuator, and c a proportionality constant.

In case of collision the translation component is reset and only the rotation component is applied. One aim in designing the actuators for this experiment was to avoid the usual Braitenberg-type, two-wheeled design which in our view restricts the freedom of evolution and introduces unnecessary
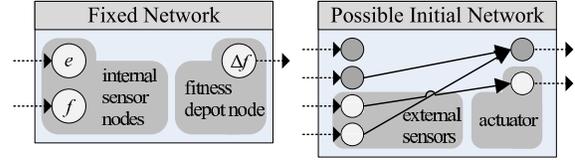
design bias by pre-specifying a bilaterally symmetrical body plan. Our actuators can be thought to act like small jets or flagella. An agent can possess any number and configuration of actuators and is therefore less restricted in the locomotion strategies it applies.

B. Neural Controller Networks

The agents are controlled by a neural network. The structure and properties of these controller networks are evolved. We use nodes with piecewise linear transfer functions where two parameters define the operating range of the neuron. Both of these parameters are initially random and evolved for each node individually. A link (synapse) is defined by a real valued weight $w \in (-\infty, \infty)$ and the two nodes it connects. Weights are at present heritable, no lifetime learning takes place. During reproduction not only parameters of existing structure can be modified, but also all components can be added or removed by mutation operators to form *completely arbitrary* (also recurrent) network structures.

First generation agents start out with a minimal network that consists of three fixed nodes and a small random number of sensors, actuators, and connections (see fig. 4), all initially set to random values.

C. Metabolism

As discussed in section II energy is the limiting resource that drives the evolution. Because energy is the only resource every aspect of an agent is somehow connected to its energy budget. Every property and every structural component has an effect on an agent's metabolism. This is necessary in this model to create an evolutionary gradient.

The metabolism of an agent has a constant and a variable component. The constant part M_{const} depends on the properties of the agent's body (i.e. its mass m) and the number of nodes $|N|$ and links $|L|$ in the neural network while the variable part consists of the actuator activation and the activation of the controller network. The relationships are given by:

$$M_{const} = \sqrt{c_1|N| + |L|} \cdot c_2 + m \cdot c_3 \qquad (7)$$

$$M_{act} = m \cdot c_4 \sum_{i}^{|A|} o_i \qquad (8)$$

$$M_{net} = c_5 \sum_{i}^{|N|} o_i \qquad (9)$$

$$M_{total} \quad = \quad M_{const} + M_{act} + M_{net} \qquad (10)$$

where (8) sums over the activation o_i of all actuators $|A|$ and (9) sums over the outputs o_i of all network nodes $|N|$. The c_i are metabolic proportionality constants.

D. Sensor-Motor Loop

The sensory-motor loop of the agents includes the following major steps at every iteration of the main program loop. After all the sensors have been updated with the current values the signals are propagated through the controller network. This results in the current action of the agent being executed after taking any collision management into account. Finally, the energy balance is drawn for the agent and the environment (see section V-C).

V. ECOSYSTEM DYNAMICS

Considerable effort has been put into providing potentially rich agent-environment interactions along with a complex agent metabolism to allow for the evolution of diverse and non-trivial survival strategies. Bearing this in mind, computational cost must always be another important consideration in simulation. Since we were primarily interested in observable agent behaviours we chose a foraging scenario with mobile artificial agents because it is a well-studied scenario where some likely behaviours can be partly anticipated and often intuitively interpreted by an observer.

A. Resource Consumption and Renewal

If an agent collides with an energy source, a certain amount of energy is transferred from source to agent. The amount of energy ΔS_E transferred from a source is determined by the maximum absorption $\Delta\epsilon(t)$ of an agent a as defined in (1) but cannot exceed the current energy content S_E of the source at the time.

$$\Delta S_E = \min(S_E, \Delta\epsilon_a(t)) \qquad (11)$$

Energy sources are therefore exhaustible but they 're-grow' at a slow constant rate μ.

$$S_E(t+1) = S_E(t) + \mu - \sum_{a \in A} \Delta S_{E,i} \qquad (12)$$

The size of an energy source equals its energy content. An energy source has a solidness of zero and allows an agent to pass through unhindered.

B. Collision

In addition to stopping the agent, an energetic cost (damage) is assigned to a collision. This cost depends on numerous properties of the agent's body as well as its speed and angle to the object and thereby 'grounds' its morphology and behaviour in the virtual world. Even though the damage calculation is very simple compared to a physically realistic model, it is differentiated well enough to allow the evolution of very different coping strategies, both on the morphological as well as the behavioural level. Especially in the case of agent-agent collisions, a damage model that is connected to the agent's morphology opens the possibility of different
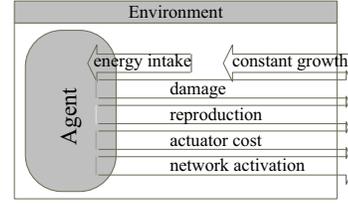

Fig. 5. Total energy balance of agents and environment.

interaction strategies. In theory even the emergent evolution of a predator-prey scenario is conceivable, if agents 'deliberately' collide with other agents in order to get their stored energy. The damage calculation is based on the length of the difference vector between the movement vector $\vec{v}_A$ of an agent A and the movement vector of the object it collides with (which, in the case of an immobile object, is the zero vector). Because the collision management resets the translation component in case of collision, we calculate two damage factors df_E and df_I. The former corresponds to the *effective* movement $\vec{v}_{A,E}$ of the agent and the latter to the *intended* movement $\vec{v}_{A,I}$. The effective movement only has an effect at the time of the actual 'crash'. The intended movement continuously affects the damage while the agent 'pushes' against an object. The two costs are balanced by proportionality constants $c_{d,E}$ and $c_{d,I}$, where $c_{d,E} > c_{d,I}$.

There is no global physics engine that calculates these costs. Instead, agents calculate their portion of a collision for both objects involved during their update cycle. The costs are given by [1]:

$$df_E \quad = \quad c_{d,E} \cdot \min(|\vec{v}_{A1,E}|, |\vec{v}_{A1,E} - \vec{v}_{A2,E}|) \quad (13)$$

$$df_I \quad = \quad c_{d,I} \cdot \min(|\vec{v}_{A1,I}|, |\vec{v}_{A1,I} - \vec{v}_{A2,I}|) \quad (14)$$

$$D_{A1 \rightarrow A1} \quad = \quad (df_{A1,E} + df_{A1,I}) \cdot \min(m_{A1}, m_{A2}) \quad (15)$$

$$D_{A1 \rightarrow A2} \quad = \quad (df_{A1,E} + df_{A1,I}) \cdot m_{A1} \qquad (16)$$

where in (15) the sum of the two damage factors is weighted by the minimum mass m_{Ai} of the two objects to yield the damage to the colliding agent and (16) describes the damage received by the passive agent. If the second agent also actively collides with the first agent (e.g. in a head on collision) it also calculates a damage.

C. Energy Balance

The metabolism of the agents together with the described dynamics determine the total energy budget of the ecosystem which is updated every timestep. The ecosystem is not a closed system with respect to energy. Energy is added from outside the system through the renewable resources described in (12) and dissipates through the metabolic consumption of the agents according to (10). A schematic illustration of the total energy balance is shown in fig. 5.

---

[1]For the sake of clarity, another factor to compensate for artifacts of discretisation and the sequential update of agents is omitted here.

TABLE I

PARAMETER SETTINGS FOR THE SIX ENVIRONMENTS INVESTIGATED.

Setup	Energy Sources	Obstacles
E1	25	25
E2	25	0
E3	50	50
E4	50	0
E5	75	75
E6	75	0

VI. SIMULATION SETUP AND RESULTS

A. Simulation Setup

In this first set of experiments, six different environments have been used to empirically explore conditions where evolution reliably occurs. Evolution does not occur if the MEAM is unable to create agents that successfully reproduce or fail to establish a sustainable population. All of these environments were relatively small square arenas (100x100 units) with varying numbers of uniformly distributed obstacles and energy sources (see table I). Energy sources were initialised with an energy capacity of 1.0 and obstacles were initialised with a solidness of 1.0. The minimum enforced number of agents was 15 in all runs. In all runs, we sampled the environment and the agent population every 1000 timesteps to get statistical data. Some of these data will be discussed in the next section.

B. Results

Evolution reliably occurred in all tested environments. To see whether or not evolution is actually happening, all agents are assigned a value we call *phylogenetic generation* (PG) which indicates the number of their direct ancestors. An agent randomly created by the MEAM has a PG of 0, its offspring 1, and so on. An exemplary graph for the development of the average PG of an agent population is shown in fig. 6. If the MEAM is unable to create any viable agents, the PG of all agents in the population is always zero and no evolution occurs. Because the viability of an agent in our model is greatly determined by its environment, continual evolution can only occur if a particular agent manages to spawn a sustained population of agents. We call such an agent a *founder*. Once an agent achieves this and the MEAM is inactive, the population quickly contains only agents which descend from the founder. Especially early in an evolutionary run, when the agent's survival and reproductive success is more governed by chance than any well-adapted strategy, extinction can occur if the whole population dies out. In this case, all 'genetic' material is lost and evolution has to start over. However, in all of the conducted experiments this happened only in very early in the evolution. Later in evolution, total extinction did not usually occur and the resulting graph for maximum PG in most runs was generally monotonically increasing. A decrease in maximum PG would signify the extinction of the 'oldest' lineage.

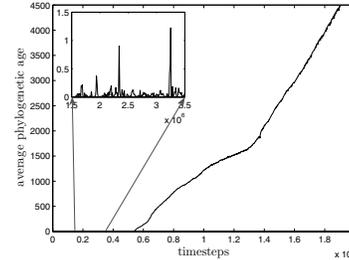The notion of phylogenetic generation in our model, however,

Fig. 6. Development of the average phylogenetic generation of a population over an exemplary run. The inset shows cases of extinction early in evolutionary time, where all inheritable material is lost.
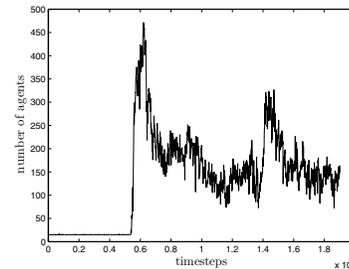

Fig. 7. The average population size in a typical run undergoes an initial phase of exponential growth once a founder agent occurs. After that it either oscillates within a certain range, or nearly stabilises.

only contains information on whether heritable information is conserved through a turnover of generations or not. It does not, in itself, give any information about the course of evolution, the level of adaptation of the agents with respect to their environment, or the complexity of their survival strategies. In a control experiment, where mutations are switched off and agents can produce only perfect clones, the PG will increase once a sustained population occurs, even though there is no evolutionary 'progress'.

Another aspect that can be observed on the population level is the average age of the current population at the time the ecosystem is sampled. As illustrated in fig. 8, before evolution occurs (at roughly $5.5 \cdot 10^7$ timesteps) the average age of the population is higher than later in the run. After the initial fast growth of the population (see fig. 7) this may be explained by the rapid decrease of free resources. However, average population age keeps decreasing even when the population size declines. Further analysis off these relationships are still required, but even so, this might hint at the idea that agents with an aggressive reproduction strategy and short average lifespans have a selective advantage. Because change and adaptation can only occur through turnover of generations, it is interesting to note that long lifespans are not necessarily selected for (as was also observed in [2]).

Due to space restrictions in this paper we will only analyse the trends in phylogenetic development for the lineages of a few exemplary individuals. A more thorough analysis of the results is certainly necessary to understand the full interaction
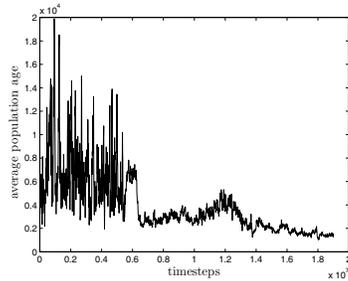
Fig. 8. The average age (in timesteps) of the agent population sampled over evolutionary time.
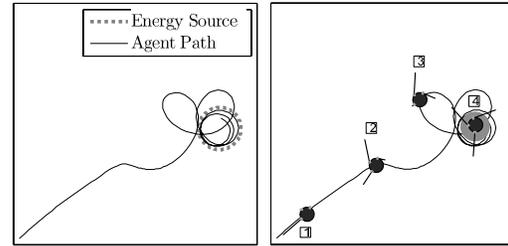

Fig. 9. Path of Agent approaching an energy source (left). The illustration on the right gives an impression of how this behaviour is achieved. (1) only the main actuator is active (the length of the line shows the activation intensity), creating a straight movement, (2) the agent senses the energy source for the first time and begins circling (3) the agent still circles but the main actuator is now completely inactive (4) the agent managed to position itself in the centre of the energy source and rotates on the spot.
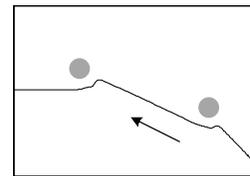

Fig. 10. Path of an agent avoiding two obstacles.

dynamics on a mechanistic level. For now we will content ourselves with looking mainly at macro-scale behaviour. Also, we do not yet have enough experimental data on the different environmental conditions described in table I to make strong claims about how the these setups influence the prevalent strategies adopted over time. However, some trends can be identified even at this early stage, and will be discussed briefly.

1) Agent-Environment Interaction: In order to successfully reproduce, agents need to develop strategies that allow them to acquire surplus energy. We will first discuss three behaviours which require locomotion and which reliably evolved in our experiments. In all experiments described, agents followed an active foraging strategy. Active foraging means that the agents moved around their environment in order to find the required energy. In contrast, a passive foraging strategy (which is theoretically possible because the energy sources are not stationary, and, indeed, occurs under certain environmental conditions not discussed in this paper) does not necessarily involve any movement on the part of the agent. In that case, agents will minimise all metabolic activity and 'wait' for the energy to come to them.

Exploration: when following an active foraging strategy, it is necessary to move around the environment even in the absence of external stimuli. We call this behaviour exploration or base movement. Founder agents typically have an actuator which is activated in the absence of external stimuli. Depending on the tilt of the actuator, early agents will usually exhibit a circular base movement of some sort. However, the evolution of exploratory behaviour was strongly influenced by the density of objects (mainly that of obstacles) in the environment. In environments with a low object density a strong trend toward faster base movement combined with a tendency to 'straighten' the forward motion could be observed. The transition from circular to straight movement was achieved either by changing the angle of the main actuator, or by adding additional actuators to counterbalance the rotational component of the main actuator. On the other hand, agents evolved in environments with a higher density of obstacles evolved much slower base movement and often did not show the tendency to straighten their base movement.

Foraging: by foraging behaviour we mean the change in agent

behaviour in the presence of energy. Early agents usually do not exhibit any foraging behaviour but have to rely on chance to pass through enough energy sources to collect enough surplus energy for reproduction. A very simple foraging strategy is to simply slow down if an energy source is sensed and in extreme cases, agents come to a complete stop. The most 'sophisticated' foraging behaviour observed so far requires the coordination of at least two (but usually three) actuators and involves 'circling in' on the energy source, and, in the most successful variants, rotating on the spot in the centre of the source. A schematic illustration of this behaviour is shown in fig. 9.

Obstacle Avoidance: the most effective obstacle avoidance behaviour (fig. 10) does not necessarily evolve in environments with the highest obstacle density. Under the conditions tested so far, the best obstacle avoidance behaviour evolved in runs with environment E1. It seems that due to the calculation of collision cost, the evolution of obstacle avoidance is closely correlated with the speed of base movement and the mass of the evolved agents. In 'lightweight' and slow-moving populations, obstacle avoidance often did not evolve at all (at the end of the run).

2) Agent-Agent Interaction: In our present model there is no direct way for an agent to discriminate another agent from an energy source or an obstacle. The only difference is that another agent provides information on all three channels. Even though this would arguably make it difficult to ascribe any sort of deliberateness to agent-agent interactions, it is still worth looking at what happens if an agent encounters another agent. However, there are at least two additional factors limiting possible variety in the evolution of agent-agent interaction.

Firstly, only actuators for locomotion are currently defined in our model and this certainly restricts the number of meaningful interactions. Secondly, in all experimental runs so far, all agents of a population were quite similar to each other and therefore, all acted similarly when encountering another agent in a certain situation.

Basically, if an agent reacts to another agent at all, it will either treat a fellow agent as an energy source, or as an obstacle. If both obstacle avoidance and foraging behaviour can be observed in an agent, it will usually also react to another agent. The reaction in that case often depends on the current level of the reproductive depot of the other agent. If this is low, the agent is likely to avoid the other agent. However, if the level of the reproductive depot of the other agent is high, the agent might approach and even chase the agent. Purely from the viewpoint of an external observer, this can lead to curious situations where an agent is furiously chased by another and, if the first agent reproduces (and therefore, its 'energy signature' drops), the chasing agent suddenly loses interest and wanders off. Note however, that since there is no clear speciation (however this might be defined in this context) there is no real predator-prey relationship, as the role of aggressor is interchangeable and determined only by the state of the two agents at the time of the encounter.

3) Reproduction Strategies: The principle behind the typical evolved strategies for reproduction is quite simple and can in most cases be summarised as follows: Invest in reproduction when energy is present, otherwise don't. However, the mechanisms of how this is achieved differ. 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 an actuator used for foraging. Many agents use a combination of these strategies. Additionally, often a negative correlation between a solidness sensor and the reproductive activity exists.

4) Morphology: As with mechanisms, we have not yet collected enough data under the various environmental conditions, to justify any strong claims about the development of different morphological aspects of the agents over the course of evolution. However, current findings suggest that the morphological development (which is of course strongly coupled with the development of certain strategies) depends on two factors.

The first factor is the initial environment. Some environments cause a strong selection pressure for certain features. For example, in environment E1 where energy and obstacles are scarce, agents always evolve a relatively fast base movement combined with small, 'lightweight' bodies. On the other hand, the slower base movement that evolves in environments with high obstacle density seems to advantage bigger and 'heavier' agents that have a higher energy capacity.

The second factor are the agents themselves. After a first burst in population growth, a lot of the environment will be other agents. Current results suggest that this can have a lasting impact on the course of evolution. However, because

the simulation becomes slow as population size increases, we currently cannot say for sure if this 'founder-bias' would eventually abate.

In general it seems that more extreme environments create larger selection pressures and cause strategies to evolve a lot faster than in more moderate environments.

VII. CONCLUSION

We have proposed a model for an evolving ecosystems approach to evolve complex agent behaviour. Evolution in our model works close to the principle of natural selection without the need for an explicit fitness function. We have introduced a novel criterion for evaluating differential reproductive success to evolve complete agents in the sense that all constituents of an agent are subject to evolution.

Our first experimental results suggest that this setup allows the evolution of complex agent behaviour even if little guidance is given to the artificial evolution process. However, we have only briefly sketched the results on an observational and behavioural level. A more thorough analysis on the functional and mechanistic level is necessary to understand the mechanisms responsible for the generation of the observed behaviour. Even though aspects of the implementation are still under development, we believe that the results suggest that our approach may be conducive for widening the scope of artificial evolution to generate agents exhibiting non-trivial behaviours in a complex ecosystem.

REFERENCES

[1] Bedau, M. A., Packard, N. H.: *Measurement of evolutionary activity, teleology, and life*. In: Langton, C. G., Taylor, C., Farmer, J. D., Rasmussen, S. (eds.): Artificial Life II. Addison-Wesley, Redwood City, CA. (1992) 431-461

[2] Menczer F., Belew, R. K.: *Latent Energy Environments*. In: Belew, R. K., Mitchell, M. (eds.): Adaptive Individuals in Evolving Populations: Models and Algorithms. Addison Wesley, Reading, MA. (1996) 191–208

[3] Menczer, F. and Belew, R. K.: *From Complex Environments to Complex Behaviors*. In: Adaptive Behavior 4. (1996) 317–363

[4] Channon, A. D., Damper, R. I.: *Evolving Novel Behaviors via Natural Selection*. In: Adami C., Belew R. K., Kitano, H., Taylor, C. (eds.): Proceedings of the 6th International Conference on Artificial Life. MIT Press, Cambridge, MA, USA, (1998) 27–29

[5] Gomez, F., Mikkulainen R.: *Incremental evolution of complex general behavior*. In: Adaptive Behavior V. MIT Press, Cambridge, MA, USA, (1997) 317–342

[6] Huelse, M., Pasemann, F.: *Modular Design of Irreducible Systems*. In: S. Nolfi, G. Baldassarre, R. Calabretta, J. C. T. Hallam, D. Marocco, J. A. Meyer, O. Miglino, D. Parisi (Eds.): From Animals to Animats 9, 9th Internaltional Conference on Simulation of Adaptive Behaviour. Springer, Berlin, Germany, (2006) 534–545

[7] Ray, T. S.: *An Approach to the Synthesis of Life*. In:C. G. Langton, C. Taylor, J. D. Farmer, S. Rasmussen (Eds.): Artificial Life II, Proceedings of the Workshop on Artificial Life. Addison Wesley, Redwood City, CA, (1992) 371–408

[8] Taylor, T.: *Creativity in Evolution: Individuals, Interactions and Environments*. P. Bentley, D. Corne (Eds.): Proceedings of the AISB'99 Symposium on Creative Evolutionary Systems, (1999).