DOI: http://dx.doi.org/10.7551/978-0-262-31709-2-ch077

Evolving gene regulatory networks controlling foraging strategies of prey and predators in an artificial ecosystem

Joachim Erdei[1], Michał Joachimczak[2] and Borys Wróbel[2,3]

[1]Department of Systems Modelling, Gdańsk University of Technology, Gdańsk, Poland
[2]Systems Modelling Laboratory, IO PAS, Sopot, Poland
[3]Evolving Systems Laboratory, Adam Mickiewicz University, Poznań, Poland
erdei@evosys.org, mjoach@evosys.org, wrobel@evosys.org

Abstract

Co-evolution of predators and prey is an example of an evolutionary arms race, leading in nature to selective pressures in positive feedback. We introduce here an artificial life ecosystem in which such positive feedback can emerge. This ecosystem consists of a 2-dimensional liquid environment and animats controlled by evolving artificial gene regulatory networks encoded in linear genomes. The genes in the genome encode chemical products which regulate other genes, sense the environment (the scent of food, prey and predators), control the animat's movement, and its foraging strategy. An animat can switch multiple times in its life between two foraging strategies (with different metabolic costs): a predator can derive food from the prey, prey just from food that diffuses in the environment. When an animat consumes enough food (or prey), it produces an offspring with a mutated genome. Mutations introduce variation into the population, and this diversity together with selective pressures leads to the evolution of control for diverse foraging strategies in an ecosystem that can support hundreds of individuals.

Introduction

The ability to prey on other organisms is a distinguishing feature of animals, and multi-level complex relationships between predators and prey are the building blocks of ecosystems. Prey-predator relationships create coupled selective pressures which can lead to evolutionary arms races between genes and lineages (species). Examples of artificial ecosystems in which such pressures exists include "sticky-feet" (Turk, 2010), in which simple multicellular organisms could both be prey and be preyed upon, and systems where separate lineages of prey and predators co-evolved, such as "Spiders" (Palmer and Chou, 2012), and "Bubbleworld" (Schmickl and Crailsheim, 2006).

In biology, prey-predator relationships evolve even between the simplest, one-celled organisms. The behaviour of these single cells is controlled by gene regulatory networks: networks in which a node represents a gene (or co-expressed genes) and edges represent regulation relationships—a directed edge from one node to another means that the product of one gene regulates the expression of another, often because this product binds physically (thanks to chemical affinity to DNA) in the vicinity of another gene. Gene products play not only regulatory roles, but also catalyse chemical reactions in the cell, form intracellular or extracellular structures, including structures necessary to sense the changes in the environment or necessary to allow for cell movement.

In our previous work we evolved artificial gene regulatory networks—using a genetic algorithm, or a novelty search algorithm (Lehman and Stanley, 2011)—to process signals (Joachimczak and Wróbel, 2010), direct multicellular development (Joachimczak and Wróbel, 2008, 2012b), and control the behaviour of unicellular (Joachimczak and Wróbel, 2009) and multicellular animats (Joachimczak and Wróbel, 2012a). Other models of artificial gene regulatory networks have been evolved to match mathematical functions (Kuo et al., 2004), evolve biological clocks (Knabe et al., 2006), study dynamics of gene expressions (Reil, 1999), and evolve robot controllers (Reil, 1999; Quick et al., 2003), also using genetic algorithms and objective fitness functions. But a genetic algorithm or a novelty search algorithm is a very imperfect model of biological evolution. In biology there is no objective fitness function—fitness corresponds to the number of offspring that is produced, and can be construed as the ability to use the resources in the environment to do so.

Because the resources of the environment are always limited, the organisms compete for them, and offspring resembles parents but also varies, natural and artificial ecosystems can be analysed from the point of view of flows of energy/matter on a short time scale and from the point of view of information on how to use the resources (evolution) on the long times scale. Artificial ecosystems are very far from capturing the complexity of matter and energy transformations in Nature or the complexity of the evolutionary process (for a review, see Dorin et al., 2008). In this paper, we present a simple system in which artificial organisms (animats) obtain matter/energy from the environment and evolve. The animats metabolise food, producing waste. Matter and energy derived from food allows them to move and to produce offspring. Animats can sense the concentra-

tion of food and waste produced by other animats, and use this information to direct movement. Offspring receives a mutated genome from the parent, allowing for evolution.

The system we present here is an extension of our previous work, in which we coupled our Gene Regulatory evolving artificial Networks (GReaNs) artificial life system with a physically plausible model of a 2-dimensional liquid environment in which thousands of animats can evolve foraging behaviour (Erdei et al., 2012). The main contribution of the work here is the introduction of a simple metabolism and a model of animats that can switch their strategy from a predator to prey and vice versa. We show here a preliminary analysis of the evolutionary trajectories and of the evolved life strategies in our artificial ecosystem.

2-dimensional liquid environment with foraging animats and diffusible substances

Simulated organisms (animats) live and evolve in a toroidal 2-dimensional liquid environment. The world contains three diffusible substances: (i) food (a source of energy) diffuses from multiple points in the environment and also from killed prey, (ii) scent of prey, and (iii) scent of predators. The last two can be seen as waste products of the prey or predator cells and allow other cells to sense them. To allow for computationally efficient yet realistic simulation of diffusion, concentration of all substances is stored using a quadtree (Finkel and Bentley, 1974) in which the root node represents the whole environment, and the other nodes – subregions of this space. Each square subregion can be divided into four smaller and equal subregions, and thus each node can have exactly four children. The depth of the tree is higher for regions where either the concentration or the gradient of concentration is high (there is a separate quadree for each substance). Places where animats are located are always represented by squares of minimum allowed size (1 length unit squared). In each subregion the chemical gradients are continuous, calculated using bilinear interpolation (Fig. 1, Gribbon and Bailey, 2004); to simplify this calculation, we permit only two kinds of neighbourhood: the neighbouring squares are of the same size or the bigger neighbour borders exactly 2 smaller squares, each one 4 times smaller in area than the bigger square.

Diffusion of substances between two adjacent squares follows the Fick's law:

$$\Delta P = \frac{c \cdot d}{D} \cdot (S_2 - S_1) \cdot \Delta t \ , \qquad (1)$$

where ΔP is the amount of food by which the concentration will increase in the next step in square 1 and decrease in square 2 (provided that the current concentration is greater in square 2 than square 1), c is the coefficient of diffusion (0.1 for food, and 0.25 for scent in the experiments described here), d is the length of the common edge, D is the distance between the centres of the squares, S_1 and S_2 are the current

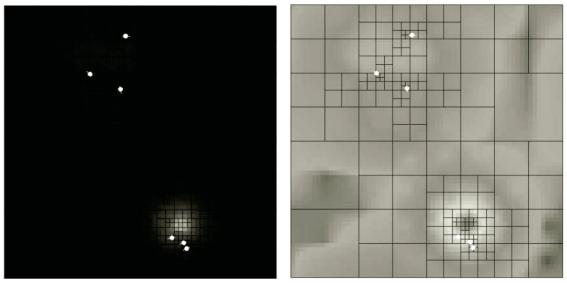

Figure 1: Modelling diffusion in 2-dimensional environment using a quadtree with continuous gradient of diffusible substances. The concentration is represented as a shade of grey, scaled linearly in the left panel and logarithmically in the right one; white corresponds to the maximum concentration close to the sources (bright squares). Animats (white circles) sense interpolated value of concentration at their exact location.

concentrations in both squares, and Δt is the duration of the simulation step. Since each square stores the concentration (not the amount), the concentration in square X is changed by $\Delta S = \frac{\Delta P}{A_X}$, where A_X is the area of square X.

At every time step, substances not only diffuse, but also degrade exponentially:

$$S(t_0 + \Delta t) = S(t_0) \cdot g^{\Delta t} \ , \qquad (2)$$

where $S(x)$ is the concentration of the substance in given square in time x, Δt is the duration of the simulation step, g is the degradation coefficient (0.99 for food, 0.9 for scent), and t_0 is the previous moment of time. Changes in concentration caused by diffusion and degradation make square areas split or merge, changing the quadtree.

One-celled predators and prey controlled by gene regulatory networks

Each animat in our system has one cell; all have equal size, shape (a circle 1 length unit in diameter), two actuators and six sensors (Fig. 2). Each sensor provides the information on the concentration of one of the following substances: food, predator scent, and prey scent (a pair of sensors for each substance) at the sensor's location. An actuator generates thrust, and because we consider our animats as models of single cells, we use a word 'flagellum' when referring to actuators (but they can be thought of as thrusters or motor-driven wheels). The activity of each actuator is controlled in a continuous fashion, so the animats can go forward or rotate in the chosen direction by varying the level of activation of the actuators. The animats move faster (accelerate by $\frac{distanceUnit}{timeStep^2}$) when both flagella are fully activated, or rotate faster (by $\frac{1rad}{timeStep^2}$) when one flagellum is fully activated and the second one is not activated at all. The maximum

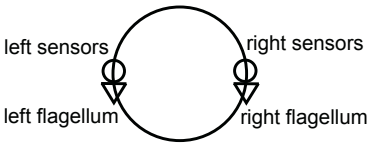

Figure 2: One-celled animat with six sensors (circles; each circle is the position of 3 sensors, one sensor for each substance in the system) and two actuators ('flagella'; triangles). The front of the animat is on the top of the figure.

speed is limited by drag (linear and angular) proportional to the velocity squared.

The concentration of food (S) in the square where the centre of the animat is located determines how much food it consumes: $0.75 \cdot S \cdot 1vu$, where $1vu$ is 1 volume unit. The food is stored internally and used up as follows:

$$\Delta M_t = (M_b + M_m \cdot \frac{a_1 + a_2}{2} + Z \cdot M_p) \cdot \Delta t \ , \qquad (3)$$

where ΔM_t is the total metabolic expenditure (by which the internal store is depleted in each time step), M_b is the base level metabolism (0.003 food units), M_m is the metabolic cost of movement (0.004), a_n is the current activation level of n^{th} flagellum (the minimal activation of a flagellum is 0, the maximum level is 1), Z is current state (prey have $Z = 0$, predators have $Z = 1$, and so does a prey cell undergoing a change into a predator; this change takes 20 simulation time steps), and M_p is metabolic cost of predation (in various experiments, we used 0.005, 0.010, or 0.015 for M_p).

An animat can choose either to feed only on the food diffusing in the environment (and run the risk of being preyed upon) or to be a predator (and have a higher cost of metabolism, which can be seen as the cost of maintaining cellular structures necessary for killing the prey and fending of other predators; in our system predators cannot feed on other predators). The switch between these two feeding strategies depends on the internal level of a chemical product encoded by the genome (Fig. 3). This concentration can change (in 0-1 range) during animat's lifetime, so multiple such strategy switches are possible (every time a threshold of 0.5 is passed).

The fact that each switch takes time (20 time steps) prevents the evolution of a strategy to change the state in response to prey or predators nearby. Depending on the state, an animat emits either a predator or prey scent. Only the collisions between predators and prey are detected, otherwise two animats in the same state (e.g., two prey cells) can overlap. Because the area of the world is large (in comparison to the average number of individuals), such overlaps happen rarely. If a prey cell is touched by a predator, the prey dies, provides 4 units of food to the predator, and whatever food was stored in its internal store to the square in which the prey cell was located.

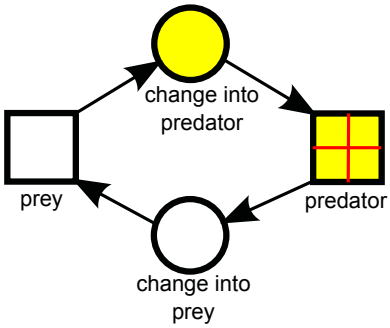

Figure 3: The possible changes between animat states. Depending on the state of its gene regulatory network, the state of the animat can change from prey (white square) to predator (yellow square; the red cross marks the state in which an organism can kill prey and cannot be killed) and vice versa, through temporary states (circles) that last 20 simulation steps. In a state marked by yellow the animat emits a predator scent, otherwise it emits prey scent.

The amount of food stored in the internal store determines if the animat is viable or can produce offspring. When the store drops to 0, the animat dies and 3 food units are released to the grid square occupied by the dead cell. When, on the other hand, there is more than 7 units of food in the store, an animat can produce one offspring cell. The cost of producing offspring is 4 units; the rest of the food in parent's store is divided equally between two cells. The new animat inherits the state of the gene regulatory network and its prey vs. predator status from the parent, it cannot change this status for the next 20 time steps to prevent a parent from immediately killing the offspring or vice versa.

Artificial gene regulatory networks that control animats' behaviour are encoded in linear genomes as described previously (Joachimczak and Wróbel, 2008). The network can be represented by a graph in which nodes correspond to chemical products in the system and the edges correspond to regulatory relations. All the products can have continuous concentrations (the minimum concentration is 0, the maximum is 1), with the exception of one special product whose level is always 1. This product serves the same role as the bias input in artificial neural network. There are 6 other input products; the concentration of these products depends on the activation of the sensors, and there are two products for each chemical substance diffusing in the environment (food, predator scent, prey scent). The concentration of one product (i_{dif}) in each pair depends on the difference in the concentration of a substance sensed by the right sensor and the concentration sensed by the left, detecting even small gradients across the body:

$$i_{dif} = \frac{1}{7.5} \cdot \log_{10}(|s_{right} - s_{left}|) + 1 \ , \qquad (4)$$

where s_{right} and s_{left} are the concentrations of the substance detected by sensors. The concentration of the second special product (i_{avg}) in each pair depends on the mean concentration sensed on the right and left:

$$i_{avg} = \frac{1}{7.5} \cdot \log_{10}(\frac{s_{right} + s_{left}}{2}) + 1 \ , \quad (5)$$

Apart from 7 input products, there are 3 output products: 2 control directly the activity of the two flagella, one determines whether an animat is currently a predator or a prey.

The topology of the regulatory network is encoded in a linear genome, which is a list of genetic elements. Each element stores 4 numbers: a type of the element, a sign, and 2 coordinates. There are 4 possible types: regulatory, coding, input, and output. Coding elements and inputs define products that have affinity to regulatory elements. A series of regulatory elements followed by a series of coding elements is a regulatory unit. An output element is a regulatory unit by itself, as if it was a regulatory unit with one regulatory element (with the sign and coordinates of the output element) and a virtual coding region coding a product that does not have affinity to any regulatory elements. Regulatory units and inputs correspond to nodes in the graph that represents the networks; the edges are defined by affinities. The affinity between two elements is determined by their coordinates. Each element defines a point in a 2-dimensional abstract space (which has nothing to do with the 2-dimensional liquid environment in which the animats move). The affinity is maximum if two points overlap, and decreases with the Euclidean distance between points, reaching zero if the distance is more than 5. If K products have affinities to the J regulatory elements of a regulatory unit, the concentrations of all products belonging to this unit L_Ω will change depending on the affinities (Euclidean distances, $d_{k,i}$) and concentrations of the regulating products in the previous step (L_k):

$$L_\Omega = \frac{2}{1 + e^{-(\sum_{j=1}^{J} \sum_{k=1}^{K} L_k (m_k \cdot m_i \cdot \frac{10 - 2d_{k,i}}{10 d_{k,i} + 1}))}} - 1 \ , \quad (6)$$

.

While the concentrations of products change during animat's life, the topology of the network (the number of nodes and edges) does not. It only changes when an offspring is produced — the parent keeps the old genome, the offspring receives a mutated copy, so genomes that encode individuals who manage to reproduce are maintained in the population (similarly to the microbial genetic algorithm; Harvey, 2011). There are 2 types of mutations that change the offspring genome: simple mutations acting at the level of a single genetic element and complex mutations, acting at the level of the whole genome. There are 3 types of simple mutations, each can occur independently, with probability 0.01: change of type, change of sign, and change of coordinates (each coordinate is modified by a random value from a normal distribution with $\mu = 0$, $\sigma^2 = 5$). Complex mutations—deletions and duplications—happen each with a probability of 0.002 per genome; the number of genetic elements removed or copied after a randomly chosen element is drawn from a geometric distribution (with a mean of 10).

Evolution of foraging strategies in an artificial ecosystem

We started each evolutionary run from 1 000 animats with randomly generated genomes consisting of 10 regulatory units (with 1 regulatory and 1 coding element in each), 3 outputs, and 7 inputs. At the start of each run, the animats had random locations in the environment 256 length units across. Each evolutionary run continued for 3 000 000 time steps, and the amount of food sources decreased linearly in time from 64 (a number high enough for the animats with random genomes to survive) to 24 at the end (increasing the selection pressure). Each food source provided 0.2 food units per time step, starting from the initial 60 food units, except for the first 64 sources, which had between 1 and 60 food units initially (a number drawn from a uniform distribution), otherwise sources would be depleted periodically. We replaced depleted source with new ones at new random locations.

We have simulated evolution for 36 independent runs in total, 12 runs for 3 values of metabolic cost of predation (M_p in Eq. 3): 0.005, 0.010, and 0.015; we kept other parameters without change (Table 1). Our previous results for the situation without predation (Erdei et al., 2012) indicate that evolution of foraging is efficient. Is it equally so in the presence of predators? How the metabolic cost of predation influences the efficiency of animats' strategies? Do animats evolve to track prey or avoid predators?

When the metabolic cost of predation was 0.015, in all the runs there were animats at the end of the run, but the lower was the cost, the higher was the chance of a population dying out (Fig. 4) with decreasing food availability over time. In particular, half of the runs with the lowest cost died out around time step 200 000, so we removed these runs from further analysis, and we show the other results only up to step 2 000 000, when the populations started to die out for the intermediate cost (all the trends we discuss, however, continue beyond this time point with no qualitative changes).

The number of individuals stabilized each time the number of food sources decreased (Fig. 5b), suggesting that the animats adapt to the new environmental conditions. When we tested the animats with random genomes in the environment with 24 food sources, they were not able to survive, indicating evolution of efficient foraging.

For the runs with the highest metabolic cost of predation, the final number of predators was small, but temporarily—when the food availability was still high—the environment supported more predators than prey (Fig. 5a). On the other hand, when the cost was 0.005 it did not pay to be prey—

Table 1: Parameters of the environment and animat metabolism in evolutionary experiments

parameter	value
food degradation coefficient	0.99
food diffusion coefficient	0.1
scent degradation coefficient	0.9
scent diffusion coefficient	0.25
duration of state change	20 time steps
no state change after division	20 time steps
base level metabolism	0.003 food units
metabolic cost of movement	0.004 food units
metabolic cost of predation	0.005, 0.010 or 0.015 food units
metabolic cost of reproduction	4 food units
reproduction threshold	7 food units
world size	256 square length units
length of evolutionary runs	3 000 000 time steps
food sources	64 to 24, decreasing with time
initial size of food sources	60 food units
source depletion rate	0.2 food units per time step

the average amount of prey in 12 runs decreased almost to 0—and when the cost was 0.010, the amount of prey was temporarily driven almost to zero, before at average it stabilized at a low level (Fig. 5a). For the intermediate cost, however, the total number of animats was the smallest (Fig. 5b), and the average amount of food stored in the animats' internal stores (amount of eaten food minus metabolic costs) per time step was the lowest (Fig. 5c), suggesting that resources were channelled to tracking prey or to avoiding predators. The amount of food stored per step was the highest when the metabolic cost of predation was the highest; because the animats avoided changing to predators, they did not suffer this higher cost at all.

We analysed in detail the behaviour of several individuals from the final populations evolved under the intermediate

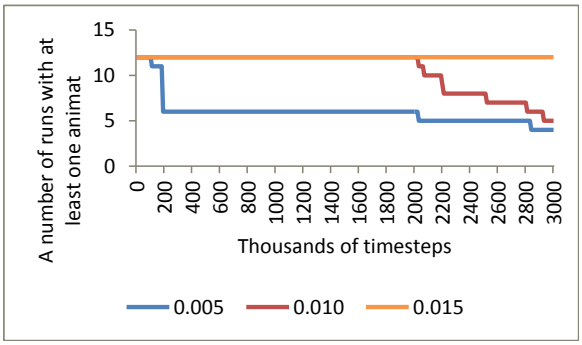

Figure 4: A number of evolutionary runs containing at least one animat as a function of time for three different values of the metabolic cost of predation.

metabolic cost of predation, by manual tracking of hand-picked individuals. This preliminary analysis suggests that a form of predator avoidance evolved in these conditions; although prey animats moved towards food sources, they avoided doing so when there were predators near the source. This strategy was efficient late in evolution because at this point prey was able to store enough food (in the internal store) to survive when searching for a food source without predators. If a prey individual moved, however, towards a source occupied by predators, it usually changed to a predator on the way. Although this change incurred a metabolic cost of producing defenses against other predators, it would pay out because high predator scent in a small area close to the source would not allow for efficient predator avoidance. On the other hand, we did not observe any predators that chased prey, even though prey did not move at full speed, and though it is easier to track another animat than to avoid it (because an animat can move faster than its scent diffuses, so the scent trail is left behind). Perhaps chasing did not evolve because the metabolic cost of movement increases linearly with speed, so a more efficient strategy is to consume food close to its source and to wait for prey there.

Conclusions and future work

We present here an artificial life system in which animats evolved life strategies that involved searching for food, avoiding predators, and waiting for prey near food sources. The animats can sense the food diffusing from sources, and the prey or predator scent diffusing from other animats. We show that it is possible to simulate the evolution of hundreds of such animats using a simple, but still realistic model of a 2-dimensional liquid environment. In particular, we modelled diffusion using a grid (represented as a quad-tree) with the resolution that adapts dynamically to concentration of chemicals and the movement of animats. The animats move in continuous space, and we approximated continuous concentration gradients using bilinear interpolation.

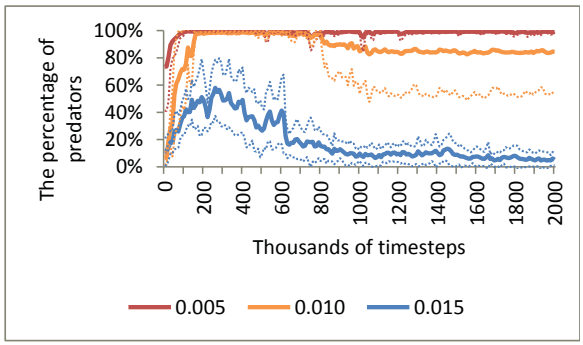

(a)

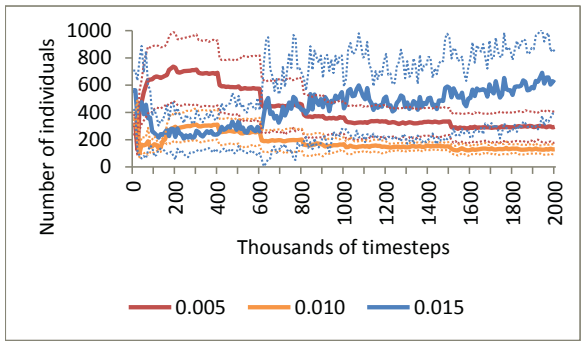

(b)

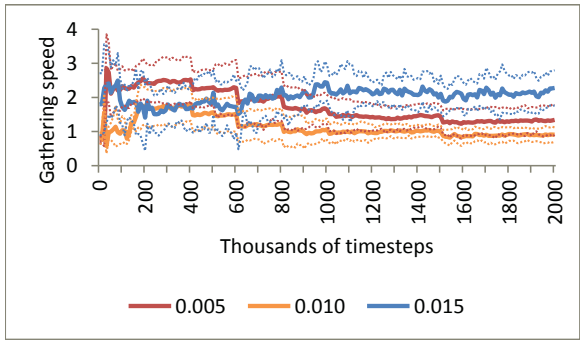

(c)

Figure 5: The percentage of predators, population size, and the amount of food stored stored by all the animats in the population per time step (gathering speed) for various metabolic costs of predation. Solid lines: averages over independent runs in which the populations did not die out before time step 2 000 000; dotted lines: averages +/- standard error.

The environmental conditions in our system can be adjusted to be qualitatively similar to those experienced by unicellular organisms. Some of such organisms—like the animats in our system—propel themselves using flagella and sense gradients across their one-celled bodies. Although many such organism live in 3-dimensional liquid environments, 2-dimensional environments (surfaces) also abound in nature, and gravitation or other forces may deliver food to such surfaces in the form of food particles. The behaviour of our animats is controlled by gene regulatory networks, and the state of this network determines if an animat is a predator (this incurs higher cost) or prey. The cost of predation can be seen as the cost of producing cellular structures necessary to kill or digest the prey, and to defend the cell against other predators. The gene regulatory networks evolve in a way that is biologically realistic, without any objective fitness function or a genetic algorithm. The survival and reproduction in our system depends on animats' ability to find food and possibly prey (in the case of predators) or avoiding predators (in the case of prey). We plan to see in our future work if more complex environments (for example, obstacles, patchiness or seasonality of the food supply) or other environmental conditions will allow in our system for the evolution of other complex behavioural strategies and to the observation of general patterns in evolution (Dorin et al., 2008) in this virtual ecosystem.

References

Dorin, A., Korb, K., and Grimm, V. (2008). Artificial-life ecosystems: what are they and what could they become? In *Artificial Life XI: Proceedings of the Eleventh International Conference on the Simulation and Synthesis of Living Systems*, pages 173–180. MIT Press, Cambridge, MA.

Erdei, J., Joachimczak, M., and Wróbel, B. (2012). Evolution of chemotaxis in single-cell artificial organisms. In *Proceedings of ICT Young 2012*, pages 185–190. Gdansk University of Technology, Poland.

Finkel, R. A. and Bentley, J. L. (1974). Quad trees a data structure for retrieval on composite keys. *Acta Informatica*, 4(1):1–9.

Gribbon, K. T. and Bailey, D. G. (2004). A novel approach to real-time bilinear interpolation. In *Proceedings of the 2nd IEEE International Workshop on Electronic Design, Test and Applications, DELTA 2004*, pages 126–134. IEEE Computer Society.

Harvey, I. (2011). The microbial genetic algorithm. In *Advances in Artificial Life: Proceedings of the 10th European Conference on Artificial Life, ECAL 2009*, volume 5778 of *Lecture Notes in Computer Science*, pages 126–133. Springer, Berlin / Heidelberg, Germany.

Joachimczak, M. and Wróbel, B. (2008). Evo-devo *in silico*: a model of a gene network regulating multicellular development in 3D space with artificial physics. In *Artificial Life XI: Proceedings of the Eleventh International Conference on the Simulation and Synthesis of Living Systems*, pages 297–304. MIT Press, Cambridge, MA.

Joachimczak, M. and Wróbel, B. (2009). Evolving gene regulatory networks for real time control of foraging behaviours. In *Artificial Life XII: Proceedings of the Twelfth International Conference on the Synthesis and Simulation of Living Systems*, pages 348–355. MIT Press, Cambridge, MA.

Joachimczak, M. and Wróbel, B. (2010). Processing signals with evolving artificial gene regulatory networks. In *Artificial Life*

XII: Proceedings of the Twelfth International Conference on the Simulation and Synthesis of Living Systems, pages 203–210. MIT Press, Cambridge, MA.

Joachimczak, M. and Wróbel, B. (2012a). Co-evolution of morphology and control of soft-bodied multicellular animats. In *Proceedings of the Fourteenth International Conference on Genetic and Evolutionary Computation*, GECCO '12, pages 561–568. ACM, New York, NY.

Joachimczak, M. and Wróbel, B. (2012b). Open ended evolution of 3d multicellular development controlled by gene regulatory networks. In *Artificial Life XIII: Proceedings of the Thirteenth International Conference on the Simulation and Synthesis of Living Systems*, pages 67–74. MIT Press, Cambridge, MA.

Knabe, J. F., Nehaniv, C. L., Schilstra, M. J., and Quick, T. (2006). Evolving biological clocks using genetic regulatory networks. In *Artificial Life X: Proceedings of the Tenth International Conference on the Simulation and Synthesis of Living Systems*, pages 15–21. MIT Press, Cambridge, MA.

Kuo, D. P., Leier, A., and Banzhaf, W. (2004). Evolving dynamics in an artificial regulatory network model. In *Proceedings of Parallel Problem Solving from Nature , PPSN VIII*, volume 3242 of *Lecture Notes in Computer Science*, pages 571–580. Springer, Berlin / Heidelberg, Germany.

Lehman, J. and Stanley, K. O. (2011). Abandoning objectives: Evolution through the search for novelty alone. *Evolutionary Computation*, 19(2):189–223.

Palmer, M. E. and Chou, A. K. (2012). An artificial visual cortex drives behavioral evolution in co-evolved predator and prey robots. In *Proceedings of the Fourteenth International Conference on Genetic and Evolutionary Computation Conference Companion*, pages 361–364. ACM, New York, NY.

Quick, T., Nehaniv, C. L., Dautenhahn, K., and Roberts, G. (2003). Evolving embodied genetic regulatory network-driven control systems. In *Advances in Artificial Life: Proceedings of the Seventh European Conference on Artificial Life, ECAL 2003*, volume 2801 of *Lecture Notes in Computer Science*, pages 266–277.

Reil, T. (1999). Dynamics of gene expression in an artificial genome implications for biological and artificial ontogeny. In *Proceedings of the 5th European Conference on Artificial Life, ECAL 1999*, volume 1674 of *Lecture Notes in Artificial Intelligence*, pages 457–466. Springer, Berlin / Heidelberg, Germany.

Schmickl, T. and Crailsheim, K. (2006). Bubbleworld.evo: Artificial evolution of behavioral decisions in a simulated predator-prey ecosystem. In *From Animals to Animats 9: Proceedings of the 9th International Conference on Simulation of Adaptive Behavior, SAB 2006*, volume 4095 of *Lecture Notes in Computer Science*, pages 594–605. Springer, Berlin / Heidelberg, Germany.

Turk, G. (2010). Sticky feet: evolution in a multi-creature physical simulation. In *Artificial Life XII: Proceedings of the Twelfth International Conference on the Simulation and Synthesis of Living Systems*, pages 496–503. MIT Press, Cambridge, MA.