DOI: http://dx.doi.org/10.7551/978-0-262-32621-6-ch017

Evolution of animats following a moving target in an artificial ecosystem

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

[1]Department of Algorithms and System 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, wrobel@evosys.org

Abstract

Many biological animals, even microscopically small, are able to track moving sources of food. In this paper, we investigate the emergence of such behavior in artificial animals (animats) in a 2-dimensional simulated liquid environment. These "predators" are controlled by evolving artificial gene regulatory networks encoded in linear genomes. The fate of the predators is determined only by their ability to gather food and reproduce—no subjective function is used to select the best individuals. Food is delivered to the environment by mobile animats who are not evolved ("prey"). Our results show promise for evolving more complex behavior relevant for nanorobotics, swarm robotics, and research on the evolution of simple cognitive abilities (minimal cognition).

Introduction

The ability to find resources (such as food) is common even among the simplest organisms. Even unicellular organisms, such as bacteria, are able to sense the concentration of various chemical substances (which can serve as food) in liquid environments, and move to a place where the concentration is higher. This phenomenon is called chemotaxis, and is an example of a minimally cognitive behavior (Beer, 1996).

Minimally cognitive behaviors attract considerable interest in Artificial Life, as a possible stepping stone towards more advanced cognitive skills (Wróbel, 2012). In many such experiments, artificial neuronal networks (ANNs) were used to control artificial organisms (animats; bio-inspired robots). In one such example of the state of the art approaches to evolution of simple behavior, virtual robots—resembling arthropods—searched for prey and avoided predators using vision (Palmer and Chou, 2012).

Yet, as we mentioned above, minimal cognition can be observed in biology even in unicellular organisms. One way in which single cells process information is with intracellular signaling networks (where nodes correspond to proteins, such as kinases, or small molecules, called transmitters; a state of the node—to concentration of, for example, a phosphorylated form of a protein, or concentration of a transmitter; and edges—to regulatory interactions). Such biological networks can be seen as computational devices, and this goes also for gene regulatory networks. Here, nodes are genes and the edges are regulatory relationships (in which a gene product regulates the activity of a gene).

Artificial gene regulatory networks (aGRNs) are a model of computation inspired by such biological networks, and have been used in experiments on the evolution of directional movement in virtual mobile robots (Quick et al., 2003; Zahadat et al., 2010). The model of aGRNs we will use in this work, called GReaNs (Wróbel and Joachimczak, 2011) has been applied before for such tasks, with a genetic algorithm as a model of evolution of either unicellular animats (Joachimczak and Wróbel, 2010; Wróbel et al., 2012), or virtual soft robots (Joachimczak and Wróbel, 2012; Joachimczak et al., 2012).

Many multi-agent simulation platforms have been developed recently. BREVE (Klein, 2003) includes rigid body simulation with friction and collision detection. It proved to be useful for simulation of flocking behavior in 3D space (Spector et al., 2003). MASON (Luke et al., 2005) has been used in diverse applications including swarm multi-agent system in continuous and discrete space. DANCE (Shapiro et al., 2005) allows to model complex physical systems including body deformation. Nonetheless, neither of these platforms allow to model diffusion with efficient storing of chemical concentrations. Therefore, we developed a simple simulated environment, where artificial organisms (animats) compete for limited resources, reproducing only when they accumulate enough of them internally.

We used this platform in our recent work (Erdei et al., 2012, 2013) where we analyzed the evolution of chemotaxis. In this previously published work, we did not use a genetic algorithm—we believe that it is an imperfect model of a biological evolution, since it requires a subjective function to evaluate fitness, an emergent property in biology. Nonetheless, when we considered evolution of behavior of animats that could prey on other animats (Erdei et al., 2013), the results were somewhat disappointing. Although the prey produced a diffusing chemical, the predators failed to evolve the ability to use the gradient of this chemical to track the prey. Rather, the predators searched for immobile food

(from which another chemical diffused), and then waited for the prey attracted by food.

In the present work we analyze if tracking of mobile prey can evolve in our artificial life platform. As previously, the animats are controlled by evolving aGRNs (using GReaNs). The task considered here—tracking a movable resource—corresponds to following a leader in swarm robotics, or chasing prey by a predator in biology. Only the predators ("followers"), however, will be evolved—the movable resource will move randomly in the experiments we will describe further on, after providing a brief recapitulation of GReaNs in the next section.

Artificial environment with diffusible substances

The model of an artificial liquid environment we will describe is an extension of the model we presented previously (Erdei et al., 2012, 2013). The model allows to simulate diffusion of substances, and movement of artificial organisms (animats) in 2-dimensional environment. The environment has no boundaries (i.e., it is a two-dimensional torus)—an animat which crosses the left border appears on the right side. While organisms occupy real-valued positions, and concentration of a substance (which can serve as food or scent) is modeled as a function of real-valued position, the concentration is stored in the software in a space-discrete data structure (quadtree; Finkel and Bentley, 1974).

This structure divides space into squares of various sizes; each square corresponds to a tree node, with the square representing the whole space at the root. Each square can be divided into 4 smaller squares (when the concentration or its gradient is high), of equal size. Thus each node can have 4 children. The leaves store the information about concentration of a substance used for modeling diffusion of providing the actual concentration to the animats (at real-valued positions, with bilinear interpolation (Gribbon and Bailey, 2004); the concentrations for squares' vertices, and centers of vertices in some situations, are calculated as an average of adjacent squares' concentrations weighted by distances to squares' centers). To simplify calculation of diffusion, we imposed a rule that each square can only border on a square of the same size, a square 4 times larger in area, or a square 4 times smaller in area (see Fig. 1).

A quadtree is a compromise, in terms efficiency and precision, between dense and sparse grids, because during simulation the squares are divided (approximating a dense grid) where the concentration (or the concentration gradient) of a substance is high. Also, the squares are divided to the smallest possible (1 square length unit in size) where animats are and where substances are released. Elsewhere, the squares remain large. The concentration of each substance is modeled separately using an individual quadtree structure. Then the concentration in each point inside a square is calculated as a bilinear interpolation using these values calculated for vertices.

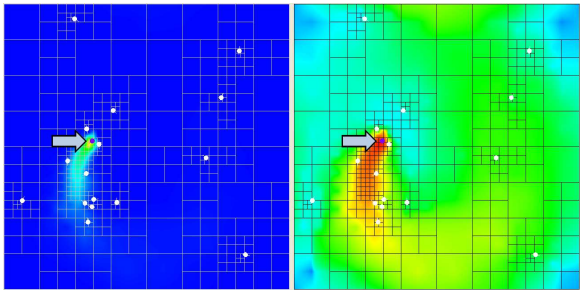

Figure 1: Modeling diffusion in 2-dimensional environment using a quadtree with continuous gradient of diffusible substances. The concentration of food is represented as hue, from blue (minimal concentration) to red (maximum concentration), scaled linearly in the left panel and logarithmically in the right one. Maximum concentration is around the prey (violet circles indicated by arrows in both panels). The predators (white circles) sense interpolated concentration values at their exact locations.

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 change in the amount of substance (by which the concentration will increase (decrease) in square 1 and decrease (increase) in square 2 if the current concentration is greater (lower) in square 2), c is the coefficient of diffusion (0.1 in the experiments described here), d is the length of the common edge, D is the distance between the centers of the squares, S_1 and S_2 are the current concentrations in both squares and Δt is the duration of the simulation step. Since each square stores concentration (not the total 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 (in the experiments described here, 0.98 or 0.97), and t_0 is the previous moment of time.

Unicellular artificial organisms (animats) controlled by aGRNs

Animats are circular, and all are 1 length unit in diameter. Each is one cell with two flagella (Fig. 2), which work like engines (i.e., generating thrust). Flagella's activations are continuous (a value between 0 and 1) so an animat can move slower or faster depending on its strategy. When both flagella are fully activated, an animat accelerates by $\frac{1du}{(\Delta t)^2}$,

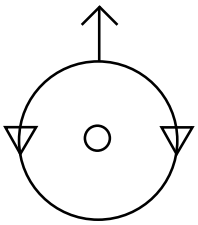

Figure 2: A unicellular animat with one sensor (circle) and two actuators ("flagella"; triangles). An arrow indicates the direction of movement when both actuators are activated equally.

where $1du$ is one distance unit, and Δt is the duration of the simulation step. When one flagellum is fully activated and the other is not activated at all, an animat accelerates its rotation by $\frac{1rad}{(\Delta t)^2}$. An animat experiences linear and angular friction which causes linear and angular drag proportional to its velocity squared. If animats collide, they bounce away from each other (in Erdei et al., 2012, 2013, collisions were not modeled).

There are two kinds of animats. One kind—the "prey"—have hand-written controllers which make them move in random trajectories. Prey do not eat or gather any substances, spend energy, die, nor reproduce. They deliver food to the environment (1 unit of food in each time step) that serves two purposes: it can be sensed by the second type of animats ("predators"), which also use it as a source of energy. In other words, prey are mobile food sources for the predators. The goal of our evolutionary experiments was to evolve predators which can follow the prey. We designed the experiments so that the only efficient way to obtain energy, necessary for reproduction, is to follow the prey. Using terminology taken from swarm robotics, the goal is to evolve animats that can "follow a leader".

In contrast to our previous work on controlling unicellular animats with aGRNs (Erdei et al., 2012, 2013), where animats had two sensors on two sides of the cells, predators in the present paper have each only one sensor for food. The sensor senses the local food concentration at the location which correspond to the cell center (Fig. 2).

A predator is controlled by an aGRN encoded in a linear genome. The model we use, called GReaNs, has been described elsewhere (e.g., Joachimczak and Wróbel 2008), and has been previously used to evolve controllers of animats (Joachimczak and Wróbel, 2010), cells in a multicellular body developing from a single cell (Joachimczak and Wróbel, 2008), and cells in a body that both develops and moves in an artificial environment (Joachimczak et al., 2012). Briefly, a genome is a list of genetic elements of 4 types: 2 types of *promoters* (*internal* and *output*), and 2 types of elements coding for *products* (*internal* and *input*). Each element is defined by 4 numbers, specifying its type, a sign, and two coordinates in two-dimensional affinity space. A series (at least one) of internal promoters followed by a series (ditto) of product-coding elements is a *regulatory unit*. An output promoter is treated as a regulatory unit with a virtual single product (not encoded in the genome explicitly and which cannot bind to promoters), whose activation level translates here to the activity of one of the animat's actuators. There are 2 possible input products: one of them (i_1) is the value received by the food sensor (s) preprocessed by the logarithmic function:

$$i_1 = \frac{1}{7.5} \cdot \log_{10}(s) + 1 \ , \tag{3}$$

and the second is always set to 1.

All products have continuous concentrations (between 0 and 1). An aGRN can be seen as a graph where vertices correspond to regulatory units, and edges correspond to product-promoter regulatory relationships. The product that regulates a regulatory unit does not need to belong to a different unit—self-regulatory relationships are permitted. When a product and a promoter have the same signs, the relationship contributes to the increase of the concentrations of all the products coded by the regulatory unit to which the promoter belongs. When the signs are different, the relationship contributes to the decrease of these concentrations. The strength of the relationship (the weight associated with the edge) depends on the Euclidean distance between the position of points (specified by the elements' coordinates) in an abstract 2-dimensional space. (This abstract space has nothing to do with the space in which the animats move, it is used solely to model chemical interactions between products and promoters.) When two points overlap, the relationship is the strongest possible. When they are further away than a prespecified distance threshold (5 units), there is no regulatory relationship. The concentrations of products coded by a regulatory unit (L_Ω) depend on their previous concentration and the concentrations of all the products with product-promoter relationships with the promoters belonging to the unit:

$$L_\Omega = \frac{2}{1 + e^{-(\sum_{j=1}^{J} \sum_{k=1}^{K} L_k (m_k \cdot m_i \cdot \frac{10 - 2d_{k,i}}{10 d_{k,i} + 1}))}} - 1 \ , \tag{4}$$

where J is the total number of promoters, K is a the total number of products and inputs in the genome, L_k is the concentration of the product k, $d_{k,i}$ is the Euclidean distance in the abstract space modeling product-promoter interactions, and m_k and m_i are element signs.

The animat gathers food ($0.75 \cdot S \cdot 1vu$ food in each time unit, where S is food concentration at animat's location, and $1vu$ is 1 volume unit), and stores it internally. This food (energy) is spent on animat metabolism and movement:

$$\Delta M_t = M_b + M_m \cdot \frac{fa_1 + fa_2}{2} \ , \tag{5}$$

where M_t is the food spent in a time step, M_b is the base metabolism cost (0.003), M_m is the movement cost (0.004), fa_1 and fa_2 are flagella activation levels (values between 0 and 1).

If all food stored internally is spent, the animat dies and decays immediately, providing 4 units food to the environment. When stored food reaches 7 units, the cell divides, and two new animats are created. Division uses 4 units of food, and the rest of food stored by the parent is divided equally between the progeny. As in our previous paper (Erdei et al., 2013), an exact copy of the parent's genome is passed to one cell, and a mutated copy to another. The inheritance without mutation by one cell is inspired by the way elitism is achieved in microbial genetic algorithms (Harvey, 2011). This cell inherits the internal product concentrations from the parent. The cell that receives a mutated genome copy starts with internal product concentrations set to zero—this is necessary because the number of the products may change after the genetic operators are applied.

The mutational process consists of two parts. First, each individual element can change its type, with probability 0.01, or its sign, also with probability 0.01. With the same probability (0.01), both coordinates of an element are modified by a random value from a normal distribution ($\mu = 0$, $\sigma^2 = 5$). Second, the element can originate (with probability 0.002) a duplication or deletion of a chunk of the genome. The number of elements copied (to a random location chosen from the uniform distribution) or deleted is sampled from the logarithmic distribution (with mean 10).

Evolution of tracking strategies

We conducted two sets of experiments, which differed in the speed of food degradation (degradation coefficients 0.98 and 0.97). The size of the environment was 256 length units across. Each experiment consisted of 7 phases. In each phase a group of predators was placed in the environment with specific number of prey (Table 1).

At the beginning of the first phase the environment was populated by 1000 predators with randomly generated genomes and 24 or 30 prey animats (Table 1). After 1 million (in the first phase) or 3 million (in all the subsequent phases) time steps, the number of prey was decreased. The prey count was high initially to enhance the survival of initially inefficient predators, and decreased afterwards so that the higher selection pressure prevents stagnation. We have conducted 20 independent evolutionary runs for each of two values of food degradation coefficient, saving the predator populations at the end of each phase (Fig. 3 and Fig. 4).

After the run ended, the saved populations were tested. In each test, only a relatively small number of predators was used because otherwise bouncing between efficient animats staying close to the prey might influence the results. To make the results as representative as possible for the actual populations, we first chose randomly 500 animats from each

Table 1: The number of prey in consecutive phases of the experiments for different food degradation coefficients.

Phase index	Number of prey for	
	coeff 0.98	coeff 0.97
1	24	30
2	16	24
3	9	20
4	7	16
5	5	12
6	3	8
7	1	4

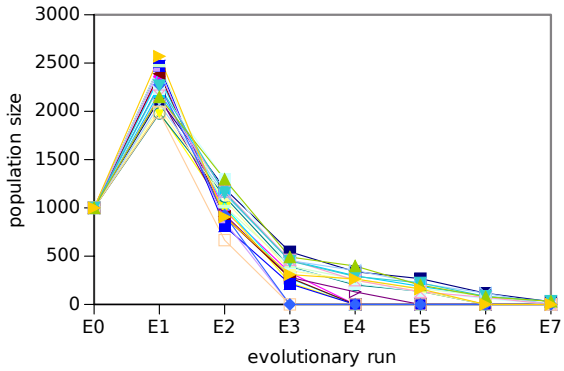

Figure 3: A size of the population for each evolutionary run after each phase (1—7; 0th phase is just 1000 animats which started the first phase) for degradation coefficient 0.98.

(with repetitions only if a population size was smaller than 500), and then separated this sample into 5 equal subsamples. These 100 predators were placed in the environment with 1 prey. During testing, the cells did not die or reproduce (and thus did not evolve). After 100 000 time steps, the mean distance at each time step between the predators and the prey was calculated, and the results for 5 subsamples were averaged. We also measured animats' progress in terms of the speed of gathering energy (the amount of eaten food decreased by metabolism costs) during the evolutionary runs. The initial population was used for comparison, but only one 100 000-step test with 100 predators and 1 prey was used for each run. In some runs no animats survived to the end of the last phase. For instance, only in 8 (for the coefficient 0.98) and 6 (for 0.97) runs there were surviving animats at the end of the penultimate phase (E6).

The results (Fig. 5, Fig. 6, Fig. 7, and Fig. 8) demonstrate that the largest evolutionary progress was made in the first 3 phases. After that the prey following behavior evolved slower or even stagnated. It is possible that the animats cannot follow the prey any closer, because there is a trade-off between moving closer and being pushed away by bouncing. Moreover, sensing changes in the direction the prey moves cannot be immediate (it requires diffusion), and we

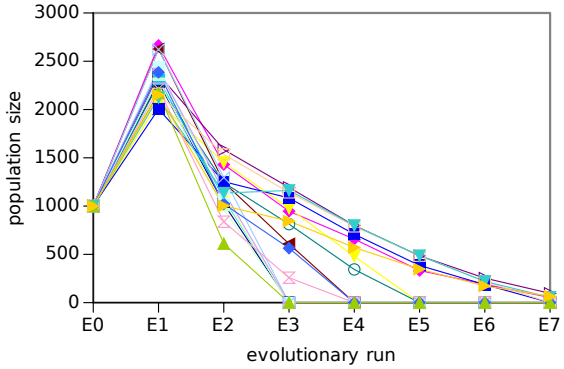

Figure 4: A size of the population for each evolutionary run after each phase (1—7; 0th phase is just 1000 animats which started the first phase) for degradation coefficient 0.97.

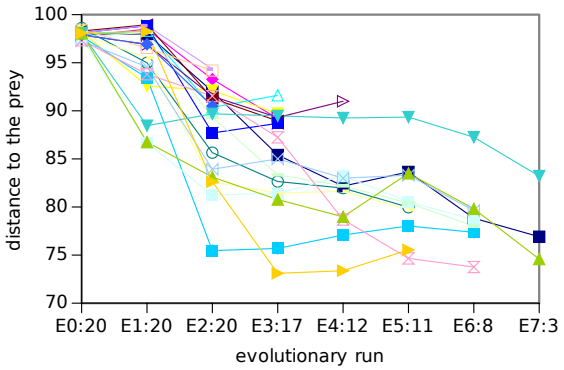

Figure 5: An average distance to the prey over 100 predators in the test environment with degradation coefficient set to 0.98. En for non-zero n is a population after n^{th} phase. $E0$ is a random population which starts the first phase. The number of populations tested (populations with survivors after each phase) is listed after a colon.

made the task even more difficult by giving the animats only one sensor. Even with those limited sensory abilities, the prey chasing behavior did evolve in the predators (Fig. 9).

Discussion and future work

We present an artificial life system where artificial organisms (animats) evolve without a genetic algorithm. The fitness is implicit and depends on the ability of an animat to find limited resources (food), necessary for reproduction with mutation. We demonstrate that in this system it is possible to evolve animats ("predators") that can follow a food source ("prey") that moves in a random trajectory, even though the predators had only one sensor of food diffusing from the prey.

The setting we used—one sensor, two actuators—is relevant for biological chemotaxis by small (not necessarily unicellular) organisms that cannot sense chemical gradients

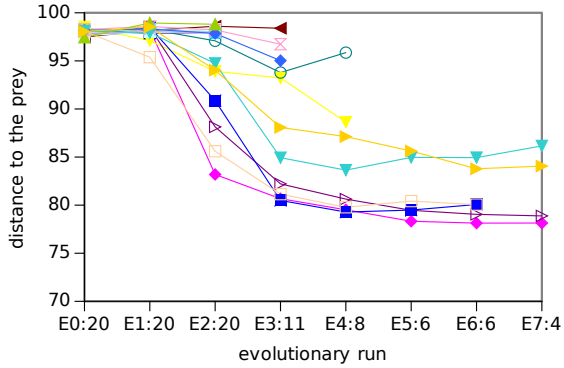

Figure 6: An average distance to the prey over 100 predators in the test environment with degradation coefficient set to 0.97. En for non-zero n is a population after n^{th} phase. $E0$ is a random population which starts the first phase. The number of populations tested (populations with survivors after each phase) is listed after a colon.

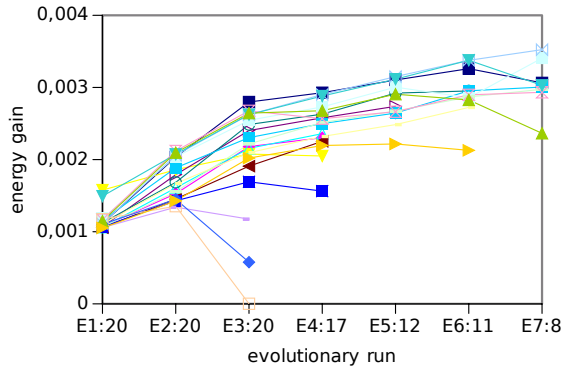

Figure 7: The average value of energy (the amount of eaten food decreased by metabolism costs) gained by one predator in one time step for each phase (E1—E7) of each evolutionary run (degradation coefficient set to 0.98). The number of populations is listed after a colon.

across their body, because their body is small relative to the differences in the concentration in their environment. Such organisms have to rely on temporal integration of information as they move. But the setting we used is not only biologically relevant—it is also applicable for cheap and small (possibly even nanoscale) mobile robots.

An even more challenging setting would be to use one sensor and one actuator; many biological organisms use actuators that can either propel the (possibly unicellular) body forward, or cause a random rotation of the body in space. This is a setting we plan to consider in the future. Even in the setting used here it will be interesting to analyze the fraction of the time in which the animat move predominately forward as opposed to local turning, in relation to the pattern in which the direction of the prey movement changes.

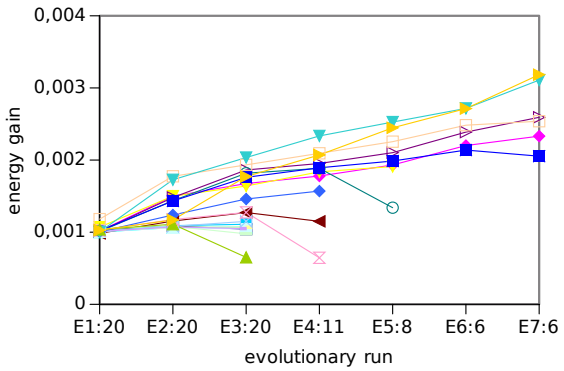

Figure 8: The average value of energy (the amount of eaten food decreased by metabolism costs) gained by one predator in one time step for each phase (E1—E7) of each evolutionary run (degradation coefficient set to 0.97). The number of populations is listed after a colon.

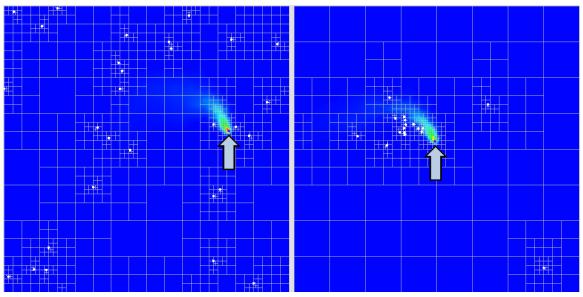

Figure 9: Prey tracking behaviour in evolved predators. The predators (white cicles) disperse all over the environment in the initial population (left), but in the runs with survivors in the last evolution phase (right), the predators remain close to the prey (violet circles indicated by arrows in both panels).

We hope that our system will allow us to test hypotheses relevant for biology and cognitive science, and evolve controllers useful in robotics. Therefore, other possible directions of future work include modifying position and number of sensors and actuators, changing how the sensors and actuators work, and introducing more challenging challenges or environments (e.g., seasonality or diversity of food sources, obstacles). This is the path towards evolution of complex cognitive abilities (Wróbel, 2012)—beyond minimal cognition displayed by predators evolved here.

Acknowledgements

This work was supported by the Polish National Science Center (BIOMERGE 2011/03/B/ST6/00399 to BW; JE acknowledges partial support of 2011/02/A/ST6/00201). We are grateful to Prof. Yaochu Jin for the discussions during BW's stay in his lab, supported by the European Commission FP7 Future and Emerging Technologies Proactive initiative: project number 270371 (COBRA). Computational resources were provided by the Tri-City Academic Computer Centre (TASK), and the Interdisciplinary Centre for Molecular and Mathematical Modeling (ICM, University of Warsaw; project G33-8).

References

Beer, R. (1996). Toward the evolution of dynamical neural networks for minimally cognitive behavior. In *From animals to animats 4: International Conference on Simulation of Adaptive Behavior*, pages 421–429.

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.

Erdei, J., Joachimczak, M., and Wróbel, B. (2013). Evolving gene regulatory networks controlling foraging strategies of prey and predators in an artificial ecosystem. In *Proceedings of the 12th European conference on Artificial Life, ECAL 2013*, pages 531–537. MIT Press.

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 *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 *ECAL 2009: Proceedings of the 10th European Conference on Artificial Life*, volume 5778 of *Lecture Notes in Computer Science*, pages 126–133. Springer, Berlin / Heidelberg.

Joachimczak, M., Kowaliw, T., Doursat, R., and Wróbel, B. (2012). Brainless bodies: Controlling the development and behavior of multicellular animats by gene regulation and diffusive signals. In *Artificial Life XIII: Proceedings of the Thirteenth International Conference on the Simulation and Synthesis of Living Systems*, pages 349–356, Cambridge, MA. MIT Press.

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. (2010). Evolving gene regulatory networks for real time control of foraging behaviours. In *Artificial Life XII: Proceedings of the Twelfth International Conference on the Simulation and Synthesis of Living Systems*, pages 348–355, Cambridge, MA. MIT Press.

Joachimczak, M. and Wróbel, B. (2012). 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, New York, NY, USA. ACM.

Klein, J. (2003). Breve: A 3d environment for the simulation of decentralized systems and artificial life. In *Proceedings of the Eighth International Conference on Artificial Life*, ICAL 2003, pages 329–334, Cambridge, MA, USA. MIT Press.

Luke, S., Cioffi-Revilla, C., Panait, L., Sullivan, K., and Balan, G. (2005). Mason: A multiagent simulation environment. *Simulation*, 81(7):517–527.

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, New York, NY, USA. ACM.

Quick, T., Nehaniv, C., Dautenhahn, K., and Roberts, G. (2003). Evolving embodied genetic regulatory network-driven control systems. *Lecture Notes in Computer Science*, 2003:266–277.

Shapiro, A., Faloutsos, P., and Ng-Thow-Hing, V. (2005). Dynamic animation and control environment. In *Proceedings of Graphics Interface 2005*, GI '05, pages 61–70, School of Computer Science, University of Waterloo, Waterloo, Ontario, Canada. Canadian Human-Computer Communications Society.

Spector, L., Klein, J., Perry, C., and Feinstein, M. (2003). Emergence of collective behavior in evolving populations of flying agents. In *Proceedings of the 2003 International Conference on Genetic and Evolutionary Computation: PartI*, GECCO'03, pages 61–73, Berlin, Heidelberg. Springer-Verlag.

Wróbel, B. (2012). Challenges for a-life approach to artificial cognition: in search for hierarchy of cognitive systems. In *Artificial Life XIII: Proceedings of the Thirteenth International Conference on the Simulation and Synthesis of Living Systems*, pages 599–600, Cambridge, MA. MIT Press.

Wróbel, B. and Joachimczak, M. (2011). Using GReaNs (genetic regulatory evolving artificial networks) for 3-dimensional asymmetrical pattern formation and morphogenesis. In *Proceedings of DevLeaNN: A Workshop on Development and Learning in Artificial Neural Networks*, pages 26–28, Paris, France.

Wróbel, B., Joachimczak, M., Montebelli, A., and Lowe, R. (2012). The search for beauty: Evolution of minimal cognition in an animat controlled by a gene regulatory network and powered by a metabolic system. In *From Animals to Animats 12: The 12th International Conference on the Simulation of Adaptive Behavior (SAB 2012)*, volume 7426 of *Lecture Notes in Computer Science*, pages 198–208, Berlin, Heidelberg. Springer Berlin Heidelberg.

Zahadat, P., Christensen, D. J., Schultz, U. P., Katebi, S., and Stoy, K. (2010). Fractal gene regulatory networks for robust locomotion control of modular robots. In *Proceedings of the 11th International Conference on Simulation of Adaptive Behavior: From Animals to Animats*, SAB'10, pages 544–554, Berlin, Heidelberg. Springer-Verlag.