

The Search for Beauty: Evolution of Minimal Cognition in an Animat Controlled by a Gene Regulatory Network and Powered by a Metabolic System

Borys Wróbel^{1,2,3}, Michał Joachimczak¹,
Alberto Montebelli⁴, and Robert Lowe⁴

¹Systems Modeling Laboratory, IO PAN, Sopot, Poland

²Evolutionary Systems Laboratory, Uniwersytet im. Adama Mickiewicza, Poznań, Poland

³Institut für Neuroinformatik, Universität & ETH Zürich, Switzerland

⁴Cognition & Interaction Lab, Höskolan i Skövde, Sweden

<http://www.evosys.org>

Abstract. We have created a model of a hybrid system in which a gene regulatory network (GRN) controls the search for resources (*fuel/food* and *water*) necessary to allow an artificial metabolic system (simulated microbial fuel cell) to produce energy. We explore the behaviour of simple animats in a two-dimensional simulated environment requiring minimal cognition. In our system control evolves in a biologically-realistic manner under tight energy constraints. We use a model of GRN in which there is no limit on the size of the network, and the concentration of regulatory substances (transcriptional factors, TFs) change in a continuous fashion. Externally driven concentrations of selected TFs provide the sensory information to the animat, while the concentration of selected internally produced TFs is interpreted as the signal for actuators. We use a genetic algorithm to obtain diverse evolved strategies in ecologically grounded animats with *motivational autonomy*, even though they lack a dedicated motivational circuit. There are three motivations (or *drives*) in the system: thirst, hunger, and reproduction. The animats need to search for *food* and *water*, but also to perform *work*. Because the value of such *work* is arbitrary (in the eye of the beholder), but affects the chances of reproduction, we suggest that the term *beauty* is more appropriate, and we name the task the *Search for Beauty*. The results obtained provide a step towards realizing a biologically realistic system with respect to: the way the control is exercised, the way it evolves, and the way the metabolism provides energy.

Keywords: minimal cognition, gene regulatory network, chemotaxis, microbial fuel cell, artificial metabolism, genetic algorithm

1 Introduction

The importance of bodily variables essential to metabolic functioning and internal and behavioural homeostatic regulation has been appreciated since the cybernetics movement of the 1940-60s (cf. [1]) and has inspired research into

self-sustaining robots required to trade off *work* with refueling needs [2, 3]) Microbial fuel cells (MFCs) comprise a form of artificial metabolism whose essential metabolic variables consist of chemical energy available to the bacteria in the anodic chamber and the hydration level of the cathode (Fig. 1, left). Bacteria convert the chemical energy into electric energy, which can be then made available to the animat. Chemical energy is provided in the form of *substrate*, which can be refined or unrefined biomass collected from the environment (cf. [4]). MFCs have already been deployed in the context of energy autonomous robotic agents ([4]), i.e. agents that are capable of refueling themselves whilst being flexible regarding the source of fuel (the *substrate*). Furthermore, work in simulation has indicated that motivationally autonomous (cf. [2]) robotic agents (MFC-powered simulated robots with minimal need-constrained action selection) may be imbued with minimal cognitive capacities such as anticipation and opportunism (e.g., [5, 6]).

In this paper we create a hybrid system in which artificial metabolism (an MFC) provides energy to the animats whose behaviour is controlled by a regulatory network inspired by the networks which are the basis of the control of all living cells. Our Artificial Life platform, GReaNs (for Gene Regulatory evolving artificial Networks) was previously used by two of us to model evolution of chemotaxis in unicellular animats [7] and evolution of soft-bodied multicellular animats [8], and was developed originally for research on artificial multicellular development [9, 10]. The model of a regulatory network in GReaNs is similar to the models used by Eggenberger Hotz [11] and other researchers (e.g., [12, 13]) in the field of Artificial Embryology. We have previously demonstrated high evolvability of GReaNs in signal processing tasks [14], and the dynamical properties of other models of gene regulatory networks (GRNs) were investigated by other authors [15, 16]. We were not the first to use GRNs to control animat behaviour (e.g., [17–19] considered wall and light following, and obstacle avoidance).

In a previous model of a hybrid (symbiotic) system – robot and MFC – proposed by two of us [20], the generated energy is stored in the capacitor (C ; Fig. 1, right), added externally to the fuel cell in order to comply with the electric constraints dictated by physical robots. Functioning of the fuel cell depends on the balance of the levels of hydration and substrate. Water and substrate need to be replenished to maintain the electrochemical process, in accordance with appropriate agent behavior. In the work reported here, animat behavior enables this replenishment, while the functioning of the MFC may mould cognitive-behavioural capacities at a level of grounding not investigated previously, i.e. using GRNs. The level of integration between the animat and the environment that GReaNs promote in combination with the energy constrained dynamics demonstrated by MFC-powered robots promises much for exploring emergent cognitive phenomena in animats.

The platform used in this paper has four biologically inspired elements: metabolism (MFC), control (GRN), evolution (a genetic algorithm), and a model of a unicellular animat which interacts with its environment in a physically-realistic fashion. This interaction is governed by simple simulated physics (which includes Newtonian laws). The way the animat senses chemicals in the environment is inspired by the mechanisms in unicellular eukaryotic organisms (for a short review, see [21]). Finally, the animats need to trade off *work* with refueling needs in order to produce progeny. Because the value of such *work* is arbitrary

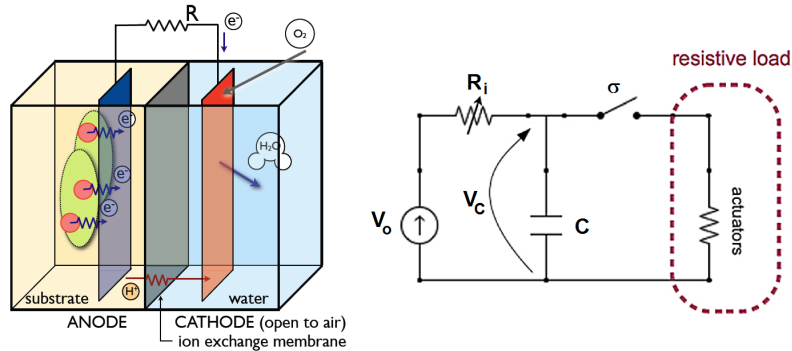


Fig. 1. The oxygen-diffusion cathode microbial fuel cell and the wiring diagram of the animat. See text for details.

(in the eye of the beholder), we suggest that *beauty* is more appropriate, and name the task the *Search for Beauty*. Another term that might capture the urge to perform *work*, with rich bio-philosophical connotations, is *striving*.

In section 2 (Model), we will first provide a brief description of the elements of the platform. In section 3 (Results), we will analyse the behaviours of animats which were evolved to perform a simple cognitive task: search for resources and *beauty*. Importantly, the information about the state of the MFC (hydration of the cathode, amount of substrate in the anodic chamber) is *not* provided to the GRN in the simulations described in this paper. Finally, in section 4, we conclude with some remarks on the implications of our results for the research in the field of adaptive behaviour.

2 Model

The model for the evolution of the linear genome encoding a GRN used in this paper is essentially the same as the one used in [7, 10, 8]. The model of the circular unicellular animat [7] with sensors at the front and two actuators at the back (Fig. 2) is modified here to include a novel energy source in GReaNs: a simulated MFC [20]. We provide here a brief description of the whole system.

2.1 Linear Genomes, Artificial Gene Regulatory Networks, and Gene Regulation

GRNs are specified by linear genomes and have *internal* and *external* nodes. A genome consists of a list of *genetic elements* (Fig. 2) of three types: *E*, *P*, or *G*. Each element is by itself a list of numbers: type (0 for *E*, 1 for *P*, or 2 for *G*), sign (-1 or 1), and two coordinates (real numbers). The genome is parsed sequentially to construct a GRN. First, all *E* elements (for *external*) are assigned to external nodes (inputs and outputs), in the order in which they appear in the genome. After *E* elements are assigned, each sequential group of *P* elements (*promoters*) followed by a sequential group of *G* elements (*genes*, which code for *transcriptional factors*, TFs) is interpreted as a regulatory unit. These units correspond to the internal nodes in the GRN. *E* elements assigned to outputs can be seen as promoters hard-wired to a product with a specific




Fig. 2. The animat and its genome. See text for details

function (controlling an actuator). Elements assigned to inputs can be seen as coding for TFs whose concentration is determined externally to the cell (e.g., by the activity of sensors). An important feature of our model is that there is no limit on the number of genetic elements in the genome, and thus on the number of nodes in the GRN, and no limit on the number of links between nodes. The number of E elements is also not limited, but only nine are assigned, two to outputs (two actuators), seven to inputs (two per each of three resources, one to a TF whose concentration is kept constant, equal to 1). Superfluous E elements in the genome are ignored.

A link between the nodes is formed if a TF has *affinity* to a promoter. Direct links between external nodes are not permitted. *Affinity* depends on the coordinates. In this paper, each element has two coordinates and thus corresponds to a point in an abstract 2D space (not to be confused with the 2D environment in which animats move). *Affinity* is determined by Euclidean distance between points (with a threshold to prevent full connectivity). The concentrations of TFs (real values in the interval $[0, 1]$) change in each simulation step. The concentrations of all TFs belonging to one regulatory unit are the same, and depend on the sum of the activation of the promoters of this unit. Activation of a promoter is the sum of the concentration of each TF that has affinity to this promoter, weighted by this affinity, taking the sign of the two elements into consideration (so regulation is inhibitory when the signs differ, and excitatory otherwise). The sum of the activation is used as an argument of a sigmoid function which produces values in the interval $(-1, 1)$. The current concentration of the TFs coded by the unit are subtracted from the value of the sigmoid function, and this end result is interpreted as the rate of synthesis/degradation to determine TF concentrations in the next step (using Euler integration). In other words, all products degrade exponentially over time unless the synthesis rate is above the intrinsic degradation rate.

2.2 Animats Powered by the Microbial Fuel Cell and Their Environment

The energy produced by the MFC and made available to the animat is a function of the stored electric potential difference across the capacitor ($\varepsilon = \frac{CV_C^2}{2}$). The voltage across the capacitor is updated with the Euler integration method in each step of the simulation of GRN activity and animat movement using the equation $\frac{dV_C}{dt} = \frac{V_o - V_C}{CR_i}$ ($C = 0.0282$). V_o , the electromotive force of the MFC, and R_i , its internal resistance, depend on the level of substrate (s) provided to the anodic chamber, and on the level of hydration (h) of the cathode: $V_o = V_{o,ref}s + V_{o,min}(1-s) - V_{o,hyd}\frac{1-h}{1-h_{lim}}$, $R_i = R_{i,ref}s + R_{i,max}(1-s) + R_{i,hyd}\frac{1-h}{1-h_{lim}}$ (voltages: $V_{o,ref} = 3.2$, $V_{o,min} = 2.8$, $V_{o,hyd} = 0.18$; resistances: $R_{i,max} = 3200$, $R_{i,ref} = 550$, $R_{i,hyd} = 600$; $h_{lim} = 0.19$ is the asymptotic hydration level, see below).

The level of substrate in the anodic chamber changes linearly with time: $s = 1 - \frac{t_s}{\tau_s}$ (t_s is the time from the last replenishment with substrate, $\tau_s = 6000$, thus the rate of substrate consumption is 10 times higher than in [20]; without this modification there is no pressure for substrate replenishment). The level of hydration is modelled as: $h = h_{lim} + h_{pos} \frac{1}{1+e^{\gamma_{pos}(t_h - \tau_{pos})}} - h_{neg} \frac{1}{1+e^{\gamma_{neg}(t_h - \tau_{neg})}}$ (t_h is the time from the last replenishment with water, $h_{pos} = 1.52$, $h_{neg} = 0.68$, $\gamma_{pos} = 0.0055$, $\gamma_{neg} = 0.031$, $\tau_{neg} = 710$, $\tau_{neg} = 600$).

There are three types of resource particles in the environment: *beauty*, *food*, and *water*. The number of *beauty* particles collected determines directly the fitness of the animat, but does not affect the MFC. There is a scalar field of scent for each resource. The scent coming from all particles of a given resource is summed. When a particle is consumed by the animat, the corresponding field changes instantaneously. If it is *water* (or *food*), t_h (or t_s) is set to 0 thus simulating the rehydration of the cathode (or the replenishment of the substrate in the anodic chamber). The scent coming from a given particle decreases with the Euclidean distance (d_{Euc}) from this particle (as $\frac{1}{1+0.2d_{Euc}}$), and reaches maximum (1) at zero distance. The activation of the sensor (S_L and S_R , Fig. 2) is equal to the value of scent at the sensor's location. In order to forage efficiently, the animat has to detect the concentration at a given location and the direction in which it changes (the gradient). The sensory information is provided to the GRN using two TFs per resource, encoded by two E genetic elements. The concentration of one TF depends on the average activation of both sensors ($\frac{2}{1+e^{-\gamma_{avg}(S_R+S_L)}} - 1$), and of one TF on the difference in their activation ($\frac{1}{1+e^{-\gamma_{dif}(S_R-S_L)}}$). In other words, the concentration of the latter TF is 0.5 when the activation of the left and right sensor is the same, and it decreases towards 0 (or increases towards 1) when the right-left difference decreases (or increases). The steepness of the sigmoid functions is set to amplify small differences or to allow for a dynamic response even when the animat is close to several particles ($\gamma_{dif} = 10$, $\gamma_{avg} = 0.5$).

The thrust forces (F_{AL} and F_{AR} , Fig. 2) generated by the actuators are proportional to the concentration of a TF associated with the corresponding output. The directions of the forces are such that when the activations of the actuators differ, the animat turns, but even when only one actuator is active, the animat moves in a loop rather than turning on the spot. When a moving animat deactivates both actuators, the motion continues due to inertia until it is brought to a stop due to fluid drag (proportional to squared velocity). This drag also imposes a maximum velocity possible.

Activation of actuators entails draining the MFC capacitor (Fig. 1). When V_c drops below a certain threshold (2.03), the switch (marked with “ σ ” in Fig. 1) opens and the distribution of energy (and thus actuation) stops while the capacitor recharges. The switch closes when V_c exceeds the upper threshold (2.90).

2.3 Evolving Efficient Controllers for Foraging in 2D Environment

Each evolutionary runs is 1000 generations of a genetic algorithm with a constant population size of 100 individuals and binary tournament selection (draw two, keep the better one). The genomes of the animats in the initial population have nine E elements and five randomly created regulatory units. Coordinates in genetic elements are randomized by drawing a random direction and a random distance from (0, 0) using a uniform distribution. Genetic operators are: changes

of element type, sign, coordinates (so that the associated point in the abstract 2D space is moved in a random direction by a distance drawn from a Gaussian distribution), and duplications and deletions of a random number of elements (drawn from a geometric distribution) at random locations in the genome. The probabilities of deletions and duplications were equal.

Before the trial, 20 *beauty*, 40 *food*, and 40 *water* particles are placed at random positions in the environment (drawn from a uniform distribution centred on the initial position of the animat) to create a *random map*. The space is open (there are no boundaries). The initial direction of the animat is random. The coordinates of the animat and food particles are represented as real numbers. The equations that govern the GRN and MFC are integrated for a specific number of time steps (7000). The genetic algorithm aims to minimize the average value of the fitness function over 10 random maps. The fitness function is: $f_{fit} = 0.8b_{dir}(1 - \frac{c_{bea}}{20})$, where (c_{bea}) is the amount of *beauty* particles collected (out of 20 on the map), while the value of b_{dir} is 1 if the animat makes at least one turn to the right and one to the left during a trial (then f_{fit} is lowered by 20%), or 1.25 otherwise (no reward). This reward promotes escaping a local optimum (a hill in the fitness landscape) which consists of moving in a loop (finding particles by chance even without any control, or with control of one actuator to tighten the loop when the scent increases [7]).

3 Results

We have performed 100 independent evolutionary runs for 1000 generations, and picked one best individual from each of the 100 final populations thus obtained. These best individuals were sorted by fitness, and 20 “best-of-the-best” were chosen for further analysis. We have re-evaluated each of these 20 individuals over 1000 random maps (instead of 10 used during the runs). The individuals collected about the same number of *beauty* particles in 7000 simulation steps (the average over 1000 trials ranged from 5.40 to 7.80, out of 20 available), differed in the strategy used to maintain the viability of the MFC, i.e., in the number of *water* and *food* particles collected (Fig. 3). The reason why this diversity is possible is that the artificial metabolism model used here (MFC) allows for some energy production when hydration is low provided the substrate level is high. Biological metabolism is similar: water is one of the products of the metabolism of carbohydrates and lipids.

To describe four of these strategies, we will use colour codes for the animats (as in Fig. 3): red for the animat that collected the most particles of all three types (*beauty*: 7.80; *water*: 10.21, *food*: 11.58), fuchsia for the second best in the search for *beauty* (7.04, *water*: 3.37, *food*: 8.75), maroon for the second best animat in *food* collected (11.08, *water*: 1.61), and green for the animat that collected the smallest average number of all particles over 1000 trials (*water*: 0.50, *food*: 6.76, *beauty*: 5.68).

Although on some maps a particular animat may be unsuccessful (Fig. 4), all four animats collected a similar average number of *beauty* particles over 1000 maps when the simulation time was extended to 21000 steps (red animat: 12.86, fuchsia: 11.95, green: 13.79, maroon: 12.68). Higher consumption of provisional resources allows the red animat to maintain a higher velocity than fuchsia, and much higher than green and maroon (Fig. 5), highlighting the fact that although




Fig. 3. Diversity of animat strategy in maintaining MFC viability. The location of the points correspond to the average number of *water* and *food* particles during 7000 time steps on 1000 random maps with 40 particles of *water*, 40 of *food*, and 20 of *beauty*.

it is possible to maintain low energy production when the cathode is dehydrated, this situation has its consequences for viability (moving too fast results in draining the capacitor and stopping, which may lead to death if energy production is too low to bring V_c over the upper threshold). It is possible to observe convergent evolution of some elements of the strategies, for example, all animats display left-right “sweeping” movements when the scent is high, this prevents missing the target by a short margin. Moreover, despite differences in strategies, one can easily observe some similarity in individual trajectories over some maps (Fig. 4).

4 Conclusions and Future Work

Our results provide a preliminary glimpse on the diversity of behaviours that evolve in agents endowed with a biologically-inspired control and an artificial metabolism, and which interact with a simple environment in a physically realistic fashion. These behaviours can be viewed as a form of *minimal cognition* [22]. The evolved animats can be considered to be energetically autonomous (they can extract from the environment the resources necessary to maintain their essential variables), *ecologically grounded*, and to have evolved limited *motivational autonomy*. The set-up explored in this paper is an example of a *three resource problem* (an extension of the *two resource problem* of [2]), in which provisional resources (here: *water* and *food*) are necessary to execute *work*. Because this *work* is not necessary for bare survival, but influences the chance of producing progeny, we name it *beauty*.

The environmental set-up for the *Search for Beauty* used in this paper provides the possibility for the evolution of three motivations: thirst, hunger, and reproduction. We use the term *motivations* in a broad, evolutionary sense, because we did not provide the animats with any motivational circuit. However, there are indirect connections between control and metabolism in our system.

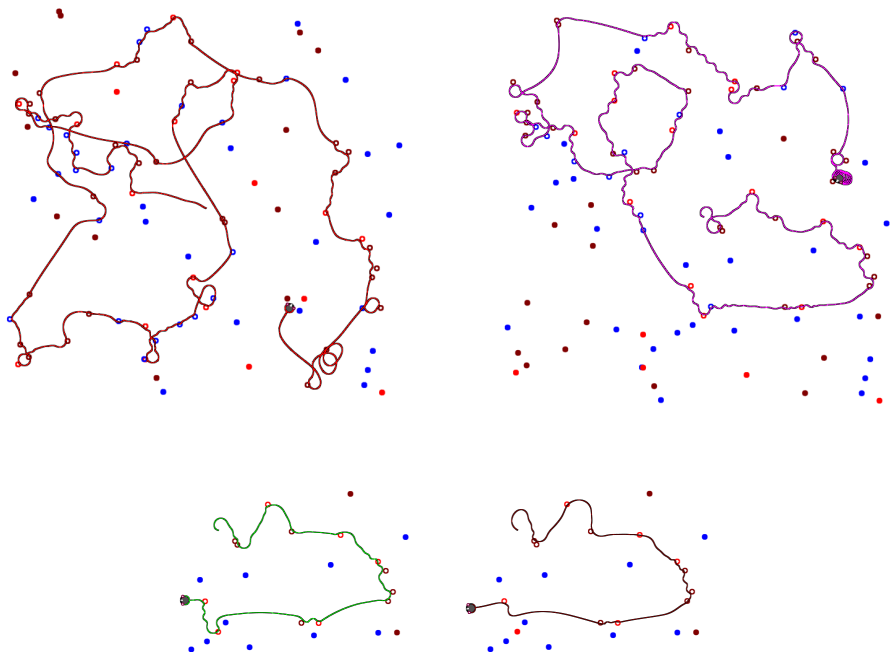


Fig. 4. Example trajectories of the evolved animats. The whole area containing particles is shown for the top two (left: red animat, right: fuchsia), only a fragment for the bottom two (left: green, right: maroon). The simulation was run until the MFC stopped its activity due to dehydration and/or lack of substrate. 40 particles of *water* (blue circles), 40 of *food* (maroon circles), and 20 of *beauty* (red circles) were placed initially, consumed particles are represented as empty circles.

On one hand, metabolism imposes constraints on the control, because if there is not enough energy to allow for movement, the animat stops (and, in principle, GRN can detect this state). On the other hand, appropriate behaviour maintains a stable metabolism. Because the connections in the GRN can be recurrent, our set-up is open to control that depends on *memory* rather than *motivation* in the strict sense of the word [23]. GReaNs can be used for tasks which require memory [14]), but the *Search for Beauty* does not: the animats do not seem to have stronger preference for *food* right after collecting *water* or vice versa (Fig. 4). Because the *Search for Beauty* can be nonetheless efficient, it is doubtful if any significant increase of efficiency in this particular environmental set-up could come from endowing animats with internal sensors of the hydration or substrate level in the MFC, or level of stored energy in the capacitor. A more promising approach is to walk along the line of further *ecological grounding*, for example, by introducing an indirect penalty for overconsumption of resources (perhaps such overeating/drinking could result in increased drag), variability in the availability of resources, or their patchy distribution [23]. Another possible means of exploring the complexity and richness of the hybrid system presented in this paper would be via providing an incentive for higher velocity, perhaps in

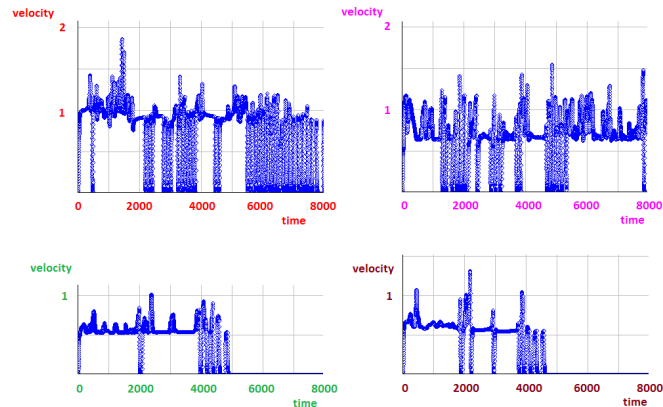


Fig. 5. Velocity of the animats. The graphs show the first 8000 simulation time steps on the same map as in Fig. 4, in the same order: red (top left) and fuchsia (top right), green (bottom left), maroon (bottom right).

a situation of direct competition for resources or in a predator/prey set-up. We plan to explore these directions in our future work.

5 Acknowledgements

Computational resources were provided by the Polish Ministry of Science and Education (project N519 384236, N303 291234), the Tri-City Academic Computer Centre (TASK), and the Interdisciplinary Centre for Molecular and Mathematical Modeling (ICM, University of Warsaw; project G33-8).

References

1. Ashby, W.R.: Principles of the self-organizing dynamic system. In: Principles of Self-Organization: Transactions of the University of Illinois Symposium. Volume 37. (1960) 125–128
2. McFarland, D., Spier, E.: Basic cycles, utility and opportunism in self-sufficient robots. *Robot Auton Syst* **20** (1997) 179–190
3. Avila-García, O., Cañamero, L.: Hormonal modulation of perception in motivation-based action selection architectures. In: From Animals to Animats 8: Proceedings of the 8th International Conference on Simulation of Adaptive Behavior, MIT Press (2004) 243–252
4. Melhuish, C., Ieropoulos, I., Greenman, J., Horsfield, I.: Energetically autonomous robots: food for thought. *Auton Robot* **21** (2006)
5. Montebelli, A., Lowe, R., Ieropoulos, I., Melhuish, C., Greenman, J., Ziemke, T.: Microbial fuel cell driven behavioral dynamics in robot simulations. In: Artificial Life XII: Proceedings of the 12th International Conference on the Simulation and Synthesis of Living Systems, MIT Press (2010) 749–756
6. Lowe, R., Montebelli, A., Ieropoulos, I., Greenman, J., Melhuish, C., Ziemke, T.: Grounding motivation in energy autonomy : A study of artificial metabolism constrained robot dynamics. In: Artificial Life XII: Proceedings of the 12th International Conference on the Simulation and Synthesis of Living Systems, MIT Press (2010) 725–732

7. Joachimczak, M., Wróbel, B.: Evolving gene regulatory networks for real time control of foraging behaviours. In: Artificial Life XII: Proceedings of the 12th International Conference on the Simulation and Synthesis of Living Systems, MIT Press (2010) 348–355
8. Joachimczak, M., Wróbel, B.: Co-evolution of morphology and control of soft-bodied multicellular animats. In: GECCO '12: Proceedings of the 14th Annual Conference on Genetic and Evolutionary Computation, ACM (2012) in press.
9. Joachimczak, M., Wróbel, B.: 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 11th International Conference on the Simulation and Synthesis of Living Systems, MIT Press (2008) 297–304
10. Joachimczak, M., Wróbel, B.: Evolution of the morphology and patterning of artificial embryos: scaling the tricolour problem to the third dimension. In: Proceedings of the 10th European Conference on Artificial Life, ECAL 2009. Volume 5777 of Lecture Notes in Computer Science. Springer (2011) 33–41
11. Eggenberger Hotz, P.: Evolving morphologies of simulated 3D organisms based on differential gene expression. In: Proceedings of the 4th European Conference on Artificial Life, ECAL 1997, MIT Press (1997) 205–213
12. Andersen, T., Newman, R., Otter, T.: Shape homeostasis in virtual embryos. *Artif Life* **15** (2009) 161–183
13. Schramm, L., Sendhoff, B.: An animat's cell doctrine. In: Advances in Artificial Life, ECAL 2011: Proceedings of the 11th European Conference on the Synthesis and Simulation of Living Systems, MIT Press (2011) 739–746
14. Joachimczak, M., Wróbel, B.: Processing signals with evolving artificial gene regulatory networks. In: Artificial Life XII: Proceedings of the 12th International Conference on the Simulation and Synthesis of Living Systems, MIT Press (2010) 203–210
15. Kuo, P.D., Leier, A., Banzhaf, W.: Evolving dynamics in an artificial regulatory network model. In: Parallel Problem Solving from Nature - PPSN VIII. Volume 3242 of Lecture Notes in Computer Science., Springer (2004) 571–580
16. Knabe, J.F., Nehaniv, C.L., Schilstra, M.J., Quick, T.: Evolving biological clocks using genetic regulatory networks. In: Artificial Life X: Proceedings of the 10th International Conference on the Simulation and Synthesis of Living Systems, MIT Press (2006) 15–21
17. Bentley, P.J.: Adaptive fractal gene regulatory networks for robot control. In: Workshop on Regeneration and Learning in Developmental Systems in the Genetic and Evolutionary Computation Conference (GECCO 2004). (2004)
18. Taylor, T.: A genetic regulatory network-inspired real-time controller for a group of underwater robots. In: Proceedings of the Eighth Conference on Intelligent Autonomous Systems (IAS-8). (2004) 403–412
19. Quick, T., Nehaniv, C.L., Dautenhahn, K., Roberts, G.: Evolving embodied genetic regulatory network-driven control systems. In: Proceedings of the 7th European Conference on Artificial Life, ECAL 2003. Volume 2801 of Lecture Notes in Computer Science. Springer (2003) 266–277
20. Montebelli, A., Lowe, R., Ziemke, T.: Towards metabolic robotics: insights from modeling embodied cognition in a bio-mechatronic symbiont. *Artif Life*, in press (2012)
21. Bagorda, A., Parent, C.A.: Eukaryotic chemotaxis at a glance. *J Cell Sci* **121** (2008) 2621–2624
22. Beer, R.: Toward the evolution of dynamical neural networks for minimally cognitive behavior. In: From animals to animats 4: International Conference on Simulation of Adaptive Behavior. (1996) 421–429
23. Saglimbeni, F., Parisi, D.: Input from the external environment and input from within the body. In: Proceedings of the 10th European Conference on Artificial

Life, ECAL 2009. Volume 5777 of Lecture Notes in Computer Science. Springer (2011) 148–155