Virtual Spatiality in Agent Controllers: A Concept to Enhance Evolvability

Jürgen Stradner[1], Heiko Hamann[1,2],
Christopher S. F. Schwarzer[3], Nico K. Michiels[3],
Thomas Schmickl[1]

[1]Artificial Life Laboratory
of the Department of Zoology,
Karl-Franzens University Graz,
Universitätsplatz 2, A-8010 Graz, Austria

[2]Department of Computer Science,
University of Paderborn,
Zukunftsmeile 1, 33102 Paderborn, Germany

[3]Institute for Evolution and Ecology,
University of Tübingen,
Auf der Morgenstelle 28, 72076 Tübingen, Germany

March 18, 2013

Abstract

Applying methods of artificial evolution to synthesize robot controllers for complex tasks is still a challenging endeavor. We report an approach which might have the potential to improve the performance of evolutionary algorithms in the context of evolutionary robotics. We apply a controller concept that is inspired by signaling networks in unicellular organisms. The implementation is based on Voronoi diagrams that describe a compartmentalization of the agent's inner body. These compartments establish a virtual embodiment, including sensors and actuators, and influence the dynamics of virtual hormones. We report results for an exploring task and an object discrimination task. These results indicate that the controller, that determines the principle hormone dynamics, can successfully be evolved in parallel with the compartmentalizations, that determine the spatial features of the sensors, actuators, and hormones.

1 Introduction

In order to control a robot or a virtual agent some device capable of processing sensor input and generating actuator output is necessary. Besides the classical way of implementing agent controllers by means of software engineering or control theory, there is research pursuing an automatic synthesis of agent controllers. Examples are the fields of evolutionary robotics [12] and adaptive agents [4]. However, the (semi-) automatic synthesis of agent controllers is still challenging, especially in complex tasks, which is partially documented by the absence of complex benchmark tasks in the literature [11].

The standard approach for controller synthesis is based on loosely biologically inspired methods, such as Artificial Neural Networks (ANN) [12]. Additionally, also the vast variety of naturally evolved control systems is often reduced to the central nervous system of vertebrates. However, the relevance of different control systems in nature is, for example, constituted in unicellular organisms that often show non-trivial behavior (without a single nerve cell), such as *Paramecium* [5]. One key feature in such chemical driven information processing system is spatiality. Nature can be said to operate on spatiality beginning with the emergence of compartments in single cells. Biological systems evolved complex membranes to establish spatiality through compartments [1] which enable eucaryotic cells to perform multiple functions that would be mutually exclusive without local separations. These functions are implemented by chemical reactions which rely on compartmentalization [9]. Additionally, conditions are established which increase the efficiency of specific tasks. This results in a strong connection between morphology and function.

The relation between function and morphology is the focus of the work presented here. In previous works, a bio-inspired approach has been proposed for controlling robots ('Artificial Homeostatic Hormone System') guided by the examples of signaling networks in unicellular organisms [13]. This paper investigates spatial effects in those controllers and the potential to boost the performance of evolutionary algorithms. We use artificial evolution and spatial partitioning by Voronoi diagrams to investigate the interaction of evolving controllers and compartments concurrently. The rationale of this approach is that, on the one hand we create an Euclidean plane of controller features such as sensor or actuator IDs and on the other hand we also embed the computational process into this space by the compartment structure and the applied diffusion processes. The evolutionary algorithm operates within Euclidean space which simplifies, for example, the encoding of partitioning sensors into two groups just by adding a line into the plane. Hence the spatial encoding might offer shortcuts during optimization. Operating within continuous space also has the advantage that mutations typically have small effects which implements an efficient local search.

2 Artificial Homeostatic Hormone Systems

We use controllers based on Artificial Homeostatic Hormone Systems (AHHS) [8, 14, 16]. AHHS are a reaction-diffusion approach. Sensory stimuli are converted into virtual hormone concentrations that possibly interact with other hormones and finally control the actuators. Hormone concentrations diffuse through a virtual inner body of the agent, and decay over time. An AHHS controller consists of descriptions of hormones (production and decay rate, diffusion coefficient) and of descriptions of rules that define the dependency between sensor input and the corresponding changes in hormone concentrations, the interactions of hormones, and the mapping of hormone concentrations to actuator control values. A rule consists of four sub-rules: sensor, linear hormone-to-hormone, nonlinear hormone-to-hormone, and actuator. The parameters of hormones and rules are encoded as arrays of floating point numbers which represent the genome. They are subject to optimization of the controller's functionality by the evolutionary algorithm. We use only one hormone and up to 30 rules; for details, see [7].

The main focus here is on spatial properties of AHHS which are defined by compartments. The compartmentalization structures the virtual inner body of the controlled agent. Hormone concentrations diffuse from one compartment to neighboring compartments as described by

$$\frac{\Delta H_h^c}{\Delta t} = D_h \nabla^2 H_h^c(t) + C, \tag{1}$$

for a constant C, that includes all other details described above, whereas the diffusion process is discrete in the implementation. The main application area of AHHS is modular robotics [8] where a natural compartmentalization due to physically connected robot modules exists. Here, we focus on internal compartmentalizations within single-module agents.

3 Compartmentalization with Voronoi diagrams

We propose an approach for evolving spatial features of controllers. This is in addition to previous works in which only functional features were adapted. We use Voronoi diagrams to describe compartmentalizations used by AHHS controllers and apply evolutionary operators similar to [15].

A Voronoi diagram is a decomposition of space, determined by the distances to a set of given points [17]. The following definition of Voronoi diagrams is based on Aurenhammer [2]. Let S be a set of points (called *sites*) in the plane. For two distinct sites $p, q \in S$ the dominance of p over q is defined by:

$$D(p,q) = \left\{ x \in \mathbb{R}^2 : |p - x| < |q - x| \right\}. \tag{2}$$

It is the subset of the plane being at least as close to p as to q. The region of a site $p \in S$ is the portion of the plane of all points dominated by p over the remaining sites as given by

$$R(p) = \bigcap_{q \in S \setminus \{p\}} D(p, q.) \tag{3}$$

All regions combined form a polygonal partition called Voronoi diagram. Informally, a region in a Voronoi diagram is all space that is closer to the site of that region than to any other sites. The Voronoi edge is the border between regions and contains points that are equally distant to both regions' sites.

Here, we use Voronoi diagrams in 2-d space to describe the compartmentalization of AHHS. Each Voronoi region corresponds to a compartment of the AHHS which holds hormone concentrations. The genome for the AHHS is extended by a Voronoi genome that consists of a set of points in 2-d coordinates which create the sites of a Voronoi diagram. We directly compute the Delaunay graph of the Voronoi sites and simulate the hormone diffusion along its connections.

Each sensor and actuator is associated with an *anchor point* in the plane, which is its virtual position. Hence, artificial evolution can be applied in two ways: On the one hand, it is an option to evolve the virtual position of the sensors and actors of the agent by mutating the *anchor points*. On the other hand, the compartment structure can be evolved. Some mutations might be silent due to their effect being too small for a change in the Delaunay graph. Others, concerning the compartment structure, are summarized in Fig. 1. Fig. 1(a) shows the original Voronoi genome to which the following three operators are applied:

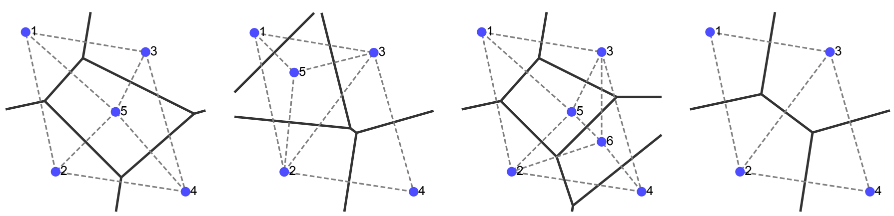

(a) Original Voronoi diagram.

(b) A shift, causing a connection to flip.

(c) Point insertion, breaking up an existing connection.

(d) Point deletion, one bridging connection is created.

Figure 1: The original Voronoi diagram shown in (a) is mutated by point shifting (b), point insertion (c), and point deletion (d). Borders of Voronoi regions are solid lines. Dashed lines are region neighborships, which is the Delaunay graph.

Point Shifting: This mutation translates a site $p = (p_1, p_2)$ to $p' = (p_1 \pm \Delta_1, p_2 \pm \Delta_2)$ for random variables Δ_1 and Δ_2. Shifting the point does trigger a connection flip. For example in Fig. 1(b), instead of the connection between points 4 and 5, there is now one between 2 and 3.

Point Insertion: When a site p with random coordinates is inserted $S' = S \cup \{p\}$, a new compartment is put into the neighborhood of cells. A point insertion can cause a change to the existing neighborhood connections, as seen in Fig. 1(c) where the new point 6 breaks up the neighborship between points 4 and 5 by being put in between.

Point Deletion: Removing a site p from the Voronoi genome $S' = S \setminus \{p\}$ causes its region in the Voronoi diagram to disappear and its space being taken over by its former neighbors. This can lead to new connections as regions that were separated before might be neighbors now. For example the regions of points 2 and 3 in Fig. 1(d) are neighbors after removing point 5.

We implemented the entire controller approach which is evolvable for functionality and spatiality into agents which had to perform two tasks of different complexity. In the first task, evolution operates on both Voronoi regions and the anchor points of sensors and actuators. In the second task, the evolution of compartmentalizations while keeping anchor points fixed is compared to the evolution of only space-independent controller features (rules and hormones). In the following these tasks are described and the results are presented.

4 Exploring task

In this scenario, the task is to explore a maze consisting of walls, see Fig. 2(a). The robot depends on its proximity sensors. The arena is divided in patches to measure the performance of the controller. Fitness is increased by visiting new patches. The maximally achievable fitness within 5000 time steps is about 5.8. A simple wall-following behavior is beneficial, but twisting trajectories might allow for traversing even more patches. The motivation to investigate this task is that the connection between the required behavior and the spatial organization of the controller is evident and we want to test whether our approach discovers this solution. A proximity has to be created between either the left actuator and the sensors pointing to the left half or the right actuator and the right-hand sensors plus an inhibitory effect of high sensor input (i.e., close wall) on the actuator. We initialize with random rules, mutate, and recombine them. We compare two variants of evolving spatiality. The first variant is based on predefined anchor points of sensors and actuators in the Voronoi plane and keeping them fixed (called 'fixed'). The second variant is to initialize these anchor points randomly and then to mutate their positions (called 'not fixed'). The population size is 200 and the number of generations is 300.

An example of an agent's behavior controlled by an AHHS with fixed anchor points presented as a trajectory is shown in Fig. 2(a). The performance comparison due to fitness between controllers using fixed anchor points in the Voronoi plane and controllers with evolved anchor points is shown in Fig. 2(d). There is no significant difference. Hence, a predefined setting is not necessary and no previous knowledge is needed in this task and effective behaviors can be evolved.

In the following, we analyze two evolved controllers: one with fixed anchor points and one with evolved anchor points. In both approaches effective controllers were evolved. We start with the controller that was evolved with fixed anchor points. The two anchor points of the actuators were placed in the upper third of the Voronoi plane, see squares in Fig. 2(b).

The anchor points of the proximity sensors are placed in one line in the lower third: triangles S_3 through S_{10} (sensors S_0, S_1, S_2 are not in use here). The evolved compartmentalization shown in Fig. 2(b) is almost symmetric and effective because it combines the sensors S_3, S_4, and S_5, that perceive the right

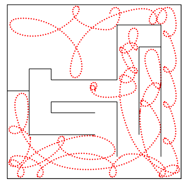

(a) trajectory, fixed

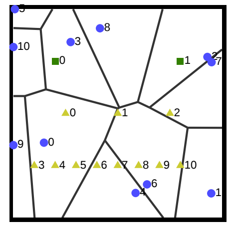

(b) Voronoi diagram, fixed

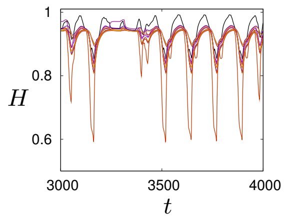

(c) hormone dynamics, fixed

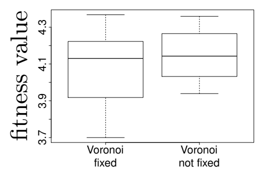

(d) comparison

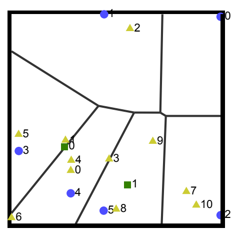

(e) Voronoi diagram, not fixed

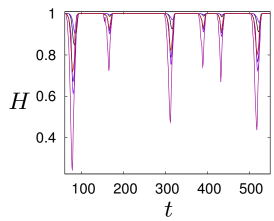

(f) horm. dynamics, not fixed

Figure 2: Analysis of evolved controllers: (a), (b) and (c) fixed sensor/actuator anchor points, fitness is 4.24; (e), (f) mutated sensor/actuator anchor points. (a) trajectory. (b), (e) Voronoi diagram, anchor points of actuators (squares), anchor points of sensors (triangles), Voronoi sites (circles). (c), (f) hormone dynamics for all compartments. (d) comparison of fixed and not fixed sensor and actuator anchor points.

side, in compartment 0 which is neighboring compartment 3 which contains the right actuator A_0. The sensors S_8, S_9, and S_{10}, that perceive the left side, are combined in compartment 6 which is neighboring compartment 2 which contains the left actuator. Summarizing the functionality of the evolved rules, the proximity sensors to the left control the left actuator inhibitorily (i.e., left wheel slows down when approaching a wall) while the right actuator always goes full speed. The trajectory contains left turns only. The hormone concentration is close to 1 in situations the agent is far from a wall and it is much lower when approaching a wall. In Fig. 2(c) drops in the hormone concentrations of all compartments are seen which are associated with approaches to walls.

The main feature of the compartmentalization for the case of evolving the anchor points as well, see Fig. 2(e), is that proximity sensor S_5 (perceiving walls to the right) and the right actuator A_0 are both in compartment 3. They are separated from the compartment containing the left actuator A_1 by one other compartment. Similarly the proximity sensors at the left side of the agent (S_7, S_8, S_9), which could disturb the sensor input from S_5, are also separated by at least one other compartment. The evolved rules are inhibitory again. The actuation of the wheels is maximal and the hormone concentration is 1 when walls are far, see Fig. 2(f). Once a wall is perceived on sensor S_5 the left actuator is not only slowed down but even turned to backward motion. This allows for a small turning radius in the right turns (trajectory not shown).

5 Object discrimination task

This scenario is an active categorical perception task, or short: object discrimination task, following [3]. Related works are [6, 10]. The robot moves in 1-d (left and right on an interval $[0, 1]$) while objects are falling down from top. The objects are either circles with different diameters or lines with different lengths. The robot has to move as close to the circles' impact positions as possible while it has to avoid those of lines. The robot has seven proximity sensors pointing to the top, distributed over an angle of $\pi/6$. At the beginning of an evaluation the robot is placed at $(0.5, 0)$. The objects' initial position is $(x, 2)$ with uniformly randomly distributed $x \in [0.1, 0.9]$ and their lengths/diameters are uniformly randomly distributed over $[0.01, 0.19]$. The robot moves only horizontally with a speed up to ± 0.05 units per step. The objects' move only vertically with constant speed of -0.03 units per step. Hence, they touch the ground after 67 steps. We have a population of 100 and 300 generations. The fitness of one evaluation is the distance d_{line} between robot and line or $1 - d_{\text{circle}}$ in the case of a circle. An controller's total fitness is an average of 60 evaluations. The motivation for this task is that the connection between required behavior and spatial controller features is not evident because it seems beneficial to consider the input of all sensors equally. Hence, we want to investigate whether a spatial approach can still be effective in this kind of tasks. We focus on evolving a) the functional features of an AHHS controller only (hormone and rule genomes), in comparison to b) additionally evolving compartmentalization. Given appropriate parameter settings of the AHHS, satisfying solutions for this task can be evolved with and without evolving compartmentalizations. In this work, we are however interested in the differences between these two approaches. Therefore, we restrict the settings to minimal AHHS controllers that are barely able to solve the task. That way we can find substantial differences in these two methods based on the maximally achieved fitness. We allow only one hormone and four rules for the AHHS controllers. We initialize the controller with one compartment. In a first set of 40 runs we allowed the evolution of compartmentalizations, hormones, and rules. In the second set of 40 runs we did not evolve compartmentalizations but only hormones and rules. The results are shown in Fig. 3(a). The 'Voronoi' approach turns out to be significantly better (Wilcoxon rank sum test, $p = 0.03947$). As stated above the system is initialized with one

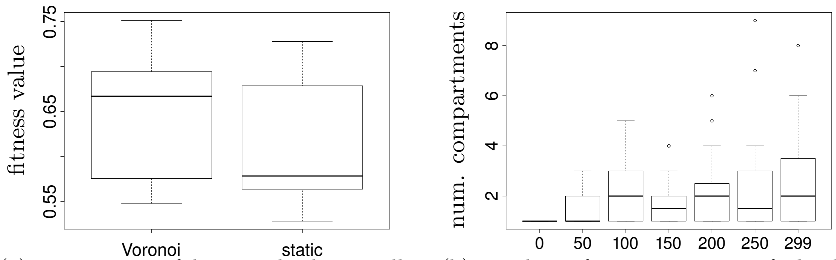

(a) comparison of best evolved controllers with and without evolving compartments based on Voronoi diagrams, $n = 40$ samples each

(b) number of compartments of the best evolved controllers, $n = 20$ samples

Figure 3: Object discrimination task, comparison of best fitness between evolving compartments and not evolving them and increase of compartments for the former case.

compartment, that is, one Voronoi point. When the compartmentalization is allowed to be evolved, new compartments are generated by adding new Voronoi points and existent points are mutated. The increase in the number of compartments is shown in Fig. 3(b). Since there is no explicit cost imposed on the number of compartments the total number is expected to increase; also because new compartments at the margins have small effects on the robot behavior.

About 65% of runs return a best controller with several compartments. In Figs. 4(a and b) we give two typical representatives of these evolved compartmentalizations. There seems to be a tendency to include the actuator in a compartment with some of the outside sensors. We can only speculate about the cause, it seems that it is relevant to react especially to objects that are not directly above the robot. The analysis of the evolved controllers is difficult in spite of the minimality of the AHHS controllers. In the following we analyze a behavior that is not prominent in its performance but in its simplicity of analysis. This controller is able to catch most of the circles that appear close to the middle and is good in avoiding lines appearing in the right half of the arena. Rule 1 mainly generates small positive actuator input for a huge hormone concentration interval $H > 0.1$. Rule 2 reduces hormone intensively for sensor input $S > 0.468$ and additionally generates big negative actuator input for $H > 0.468$. Rule 3 is mostly ineffective because it only produces small actuator input for $H \approx 0.5$ which occurs only shortly in typical runs of this controller. Finally, rule 4 is ineffective because its trigger window width is set to 0. The resulting behavior is seen in Fig. 5. This controller pursues a simple strategy of staying centered for big sensor input ($S > 0.468$) and moving to the left end of the arena if the sensor input was not big enough until $t \approx 35$. Due to the extension in the y-dimension of circles this seems to be a possible strategy which, however, has a quite high ratio of mismatches. In Fig. 5(b) the hormone dynamics for two situations, a circle or a line at position 0.45, is given. The horizontal line gives the threshold of rule 2 (0.468). In case of the circle, sensor input is big

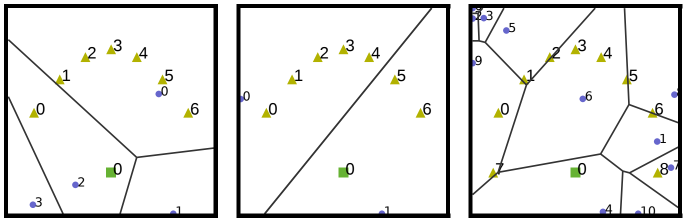

(a) standard task (b) standard task (c) noisy sensors

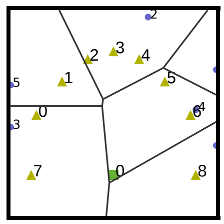

(d) noisy sensors

Figure 4: Typical evolved compartments of the best controllers for the standard object discrimination task and with two additional, disturbing noisy sensors (S_7 and S_8).

($S > 0.468$) before $t = 35$ and the hormone concentration in both compartments is reduced to zero which keeps the robot moving very slowly (due to rule 1). In case of the line, sensor input stays smaller ($S < 0.468$) before $t = 35$, the hormone concentration increases, when $H > 0.468$ the robot starts to move fast to the left due to rule 2.

In order to document this system's capabilities of separating sensors and actuators from each other by compartmentalizations we have done additional test runs. For this we have extended the object discrimination task by adding two more sensors: S_7 and S_8. We have placed them in the lower third of the Voronoi diagram to the left and the right of the actuator. These two new sensors, however, are of no purpose because they only exhibit random uniform noise over the full interval $[0, 1]$. When evolving controllers for this configuration we would expect that it is beneficial to separate sensors S_7 and S_8 from the actuator by introducing appropriate compartments. Indeed, this was observed in several of our test runs, see Figs. 4(c and d). In the case of the example shown in Fig. 4(c) most of the compartments containing sensors are directly neighboring the actuator's compartments. Sensor S_7 is as close to the actuator as several purposeful sensors. Still, by combining sensors S_2 through S_5 in one compartment might help to have a bigger effect on the actuator than sensor S_7. Sensor S_8 is separated effectively with a distance of two compartments. The example shown in Fig. 4(d) is better in separating the noisy sensors by pooling two purposeful sensors (S_5 and S_6) with the actuator in one compartment. In addition, a low diffusion rate minimizes the disturbance by the noisy sensors.

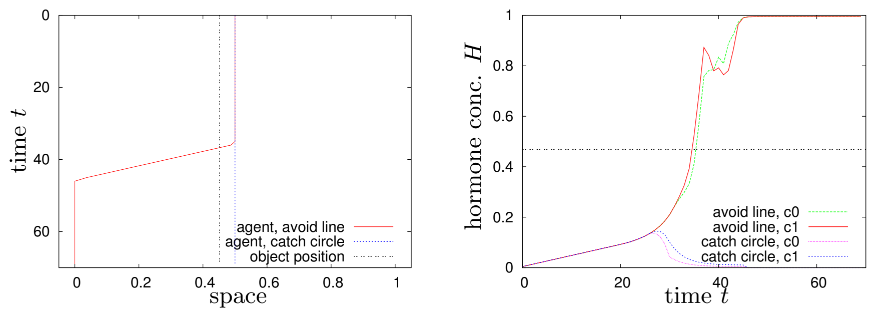

(a) agent trajectories (b) hormone dynamics

Figure 5: Agent trajectories and hormone dynamics of the analyzed controller for the two cases (avoid line and catch circle).

6 Conclusion

We have presented a concept of artificial spatiality for robot controllers inspired by compartmentalizations in natural cells. In the first task the performance of two aspects of spatiality were compared: the topology of the compartments vs. the virtual positions of the agent's sensors and actuators. This is indicated by evolving useful, symmetric compartments in case of fixed sensor/actuator anchor points, see Fig. 2(b) and especially in the case of evolved sensor/actuator anchor points where functionally associated actuators and sensors were put close together and separated from possibly perturbing sensors, see Fig. 2(e).

This is comparable to natural systems that learn while the position of their actuators change, for example, due to growing processes. In contrast to the optimization process in living systems, in artificial systems the evolutionary change of the inner morphology (i.e., the controller topology) can serve as a seed to optimize functionality. The combined evolution of morphology and function allows for alternative pathways through search space to desired agent behaviors that seem to improve the overall performance of the evolutionary approach. For example, the separation of sensors into two groups is easily achievable by the Voronoi approach (e.g., as seen in Fig. 4(b)) while it is more difficult to be evolved in a more abstract search space based, for example, on sensor IDs. An additional advantage of evolving compartmentalizations is the easy visualization and, thus, the possibility of understanding the controller's mode of operation intuitively compared to sometimes complex networks of causality in the evolved logic of AHHS rules.

The idea to apply spatial features to information processing is rather new in artificial agents. Our approach is based on concepts which are well known from nature and well tested by natural evolution. Here we were able to show promising results based on a spatial approach in autonomous agents. In future work, we plan to show the advantage of spatiality in controllers for complex tasks.

References

[1] Alberts, B.: Molecular biology of the cell. Garland Pub. (1989)

[2] Aurenhammer, F.: Voronoi diagrams — a survey of a fundamental geometric data structure. ACM Computing Surveys 23(3), 345–405 (1991)

[3] Beer, R.D.: The dynamics of active categorical perception in an evolved model agent. Adaptive Behavior 11(4), 209–243 (2003)

[4] Beer, R.D., Gallagher, J.C.: Evolving dynamical neural networks for adaptive behavior. Adaptive Behavior 1(1), 91–122 (1992)

[5] Bray, D.: Wetware: A Computer in Every Living Cell. Yale University Press (2009)

[6] Dale, K., Husbands, P.: The evolution of reaction-diffusion controllers for minimally cognitive agents. Artificial Life 16(1), 1–19 (2010)

[7] Hamann, H., Schmickl, T., Crailsheim, K.: A hormone-based controller for evaluation-minimal evolution in decentrally controlled systems. Artificial Life 18(2), 165–198 (2012)

[8] Hamann, H., Stradner, J., Schmickl, T., Crailsheim, K.: Artificial hormone reaction networks: Towards higher evolvability in evolutionary multi-modular robotics. In: Fellermann, H. et al. (eds.) Proc. ALife XII Conf. pp. 773–780. MIT Press (2010)

[9] Lodish, H., Berk, A., Zipursky, L.S., Matsudaira, P., Baltimore, D., Darnell, J.E.: Molecular Cell Biology. W.H. Freeman, New York, USA, fifth edn. (2003)

[10] Moioli, R., Vargas, P.A., Husbands, P.: Exploring the kuramoto model of coupled oscillators in minimally cognitive evolutionary robotics tasks. In: WCCI 2010 IEEE World Congress on Computational Intelligence - CEC IEEE. pp. 2483–2490 (2010)

[11] Nelson, A.L., Barlow, G.J., Doitsidis, L.: Fitness functions in evolutionary robotics: A survey and analysis. Robotics and Autonomous Systems 57, 345–370 (2009)

[12] Nolfi, S., Floreano, D.: Evolutionary Robotics: The Biology, Intelligence, and Technology of Self-Organizing Machines. MIT Press (2000)

[13] Schmickl, T., Hamann, H., Crailsheim, K.: Modelling a hormone-inspired controller for individual- and multi-modular robotic systems. Mathematical and Computer Modelling of Dynamical Systems 17(3), 221–242 (2011)

[14] Schmickl, T., Hamann, H., Stradner, J., Crailsheim, K.: Hormone-based control for multi-modular robotics. In: Levi, P., Kernbach, S. (eds.) Symbiotic Multi-Robot Organisms: Reliability, Adaptability, Evolution, pp. 240–263. Springer (2010)

[15] Schoenauer, M., Kallel, L., Jouve, F.: Mechanical inclusions identification by evolutionary computation. (1996)

[16] Stradner, J., Hamann, H., Schmickl, T., Crailsheim, K.: Analysis and implementation of an artificial homeostatic hormone system: A first case study in robotic hardware. In: The 2009 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS'09). pp. 595–600. IEEE Press (2009)

[17] Voronoi, G.: Nouvelles applications des paramétres continus à la théorie des formes quadratiques. J. für die Reine und Angewandte Mathematik 133, 97–178 (1907)