

Artificial Hormone Reaction Networks: Towards Higher Evolvability in Evolutionary Multi-Modular Robotics

Heiko Hamann, Jürgen Stradner, Thomas Schmickl, Karl Crailsheim

Artificial Life Lab of the Department of Zoology, Karl-Franzens University Graz, Universitätsplatz 2, A-8010 Graz, Austria
heiko.hamann@uni-graz.at

Abstract

The semi-automatic or automatic synthesis of robot controller software is both desirable and challenging. Synthesis of rather simple behaviors such as collision avoidance by applying artificial evolution has been shown multiple times. However, the difficulty of this synthesis increases heavily with increasing complexity of the task that should be performed by the robot. We try to tackle this problem of complexity with Artificial Homeostatic Hormone Systems (AHHS), which provide both intrinsic, homeostatic processes and (transient) intrinsic, variant behavior. By using AHHS the need for pre-defined controller topologies or information about the field of application is minimized. We investigate how the principal design of the controller and the hormone network size affects the overall performance of the artificial evolution (i.e., evolvability). This is done by comparing two variants of AHHS that show different effects when mutated. We evolve a controller for a robot built from five autonomous, cooperating modules. The desired behavior is a form of gait resulting in fast locomotion by using the modules' main hinges.

Introduction

The (semi-)automatic synthesis of robot controllers with artificial evolution belongs to the software section of evolutionary robotics (Cliff et al., 1993). The main challenge in this field is the curse of complexity because an increase in the difficulty of the desired behavior results in a significantly super-linear increase in the complexity of its evolution. This is partially documented by the absence of complex tasks in the literature (Nelson et al., 2009). Additionally, in evolutionary robotics the cost of the fitness evaluation is rather high even in case of simulations, if the application of a physics engine (simulation of friction, inertia etc.) cannot be avoided. Another challenge is the appropriate choice of a genetic encoding (Mataric and Cliff, 1996) and the basic principle of the controller design as they define the designable fraction of the search space and the fitness landscape (non-designable fractions are induced, for example, by the environment or the task itself). While the search space should be kept small, the fitness landscape should be smooth with a minimum number of local optima. Experience shows that these two criteria are contradicting. We

summarize this complex of challenges by the aim to 'strive for high evolvability'.

Concerning the problem of finding appropriate controller designs a pleasant trend can be observed in recent literature. The most prominent candidate is presumably the HyperNEAT design (Stanley et al., 2009; Clune et al., 2009). It is based on artificial neural networks (ANN) but combines the 'search for appropriate network weights with complexification of the network structure' (Stanley and Miikkulainen, 2004) through the generation of connectivity patterns. It has proven to have good evolvability combined with an adequate range of applications. Other promising, recent approaches tend to be more inspired by biology, in particular by unicellular organisms and endocrine systems. Examples showing good evolvability are the reaction-diffusion controller by Dale and Husbands (2010) and homeostasis and hormone systems based on GasNets (Vargas et al., 2009) and ANNs (Neal and Timmis, 2003). They indicate homeostasis as a prominent feature in successful adaptation to dynamic environments.

In this paper, we test and analyze a controller design called Artificial Homeostatic Hormone Systems (AHHS) that is based on hormones only and that was introduced before (Hamann et al., 2010; Schmickl et al., 2010; Schmickl and Crailsheim, 2009; Stradner et al., 2010, 2009). AHHS is a reaction-diffusion approach. Sensory stimuli are converted into hormone secretions that, in turn, control the actuators. In addition, hormones interact linearly and non-linearly comparable to the hidden layer of ANN. The topology of this hormone-reaction network is not pre-defined. Such systems show homeostatic processes because they typically converge to trivial equilibria for constant sensor input. The sensory stimuli are basically integrated in form of hormone concentrations (a form of memory) and decomposed over time (oblivion). However, during a limited period of time (transient) after a stimulus they show also variant behavior, especially, if non-linear hormone-to-hormone interactions are applied. This way, explorative behavior of the robot is implemented that allows for the testing of many sensory-motor configurations.

The desired main application of AHHS is multi-modular robotics (SYMBRION, 2010; REPLICATOR, 2010). In this field, autonomous robotic modules are studied, that are able to physically connect to each other, and can also establish a communication and energy connection. Hence, they form a super-robot called ‘organism’, that is able to re-configure its body shape, see for example, Shen et al. (2006) or Murata et al. (2008). Therefore, the underlying idea of diffusion in our reaction-diffusion system is that hormones diffuse from robot module to robot module and establish a low-level communication. Following our maxim of trying to reach a maximum of plasticity we use identical controllers in each module independent of their position within the robot organism, so there is neither a controller nor a module specialization. This concept implements the focus of evolutionary robotics on modularity (among others) in terms of hardware and software (Nolfi and Floreano, 2004). Although we evolve cooperative behaviors by evolving a kind of self-organized role selection, there is no co-evolution.

In general, our approach is more organic in contrast to the typical symbolic approach (direct encoding of pitch, roll, yaw angles, use of pattern generators using Gaussian functions etc.). The biological inspiration is not practiced as an end in itself but rather introduces more robustness in computations and it allows the diffusion of such values from module to module (implementing implicit communication).

One focus of our current research track is to design fitness landscapes by using appropriate controller designs. We investigate possibilities of smoothing the fitness landscape by a sophisticated interaction between the controller design and the mutation operator. We test whether it is useful to maximize the causality of the mutation operator (i.e., small causes have small effects) by reducing the maximal impact to the organism’s behavior. However, whether high causality is really desirable, is questionable, for example, cf. Chouard (2010).

In this paper, the investigated scenario is a modular-robotics variant of gait learning in simulation. Initially, we connect five modules in a simple chain formation as the body formation itself is not yet in our focus. The task is to move as far as possible by utilizing the hinge in each module only (no wheels).

Artificial Homeostatic Hormone Systems

In AHHS, sensors trigger hormone secretions, which increase hormone concentrations in the robot. These hormones diffuse, integrate, decay, interact and finally, affect actuators. We have analyzed AHHS controllers in single robots before (Schmickl et al., 2010; Schmickl and Crailsheim, 2009; Stradner et al., 2010, 2009). In these cases, the robot’s body was virtually divided into compartments that hold hormones and between which hormones diffuse. These compartments create a spatial context (embodiment) by associating sensors and actuators




Figure 1: Sketch of the hormone dynamics and diffusion processes in an organism. Each module holds different hormones with different concentrations, hormones diffuse through the organism based on a diffusion coefficient evolved individually for each hormone, module locations (e.g., elevation) are not relevant for diffusion. Sensor settings are simplified, actually every module has four proximity sensors.

with explicit compartments (e.g., left proximity sensor and left wheel actuator are associated with the left compartment and hence depend only on hormone concentrations of this compartment). In the case of modular robotics, the subdivision of the robot organism is naturally defined by the modules themselves. A virtual compartmentalization is not necessary and hormones diffuse from module to module (see Fig. 1). A first small case study with organisms built from three modules was reported in (Hamann et al., 2010).

AHHS1

The AHHS, as initially presented in (Schmickl and Crailsheim, 2009; Schmickl et al., 2010), is called AHHS1 in this paper. An AHHS consists of a set of hormones and a set of rules. On the one hand, it defines production/decay rates and diffusion coefficients for each hormone. On the other hand, it defines by rules the production through sensors and interaction of hormones as well as their influence on actuators. There are four types of rules. Sensor rules define the production of hormone through sensor input. Actuator rules define the control of actuators through hormone concentrations. Hormone rules define the interaction between hormones, that is, one hormone triggers the production of another hormone (or itself). Additionally, there is an idle rule to allow a direct deactivation of rules through mutations. Rules are triggered at runtime, if a certain threshold is reached (sensor values in case of sensor rules or hormone concentrations in case of hormone rules). The amount of produced hormone or the actuator control value are linearly depending on the controlling sensor or hormone respectively ($\lambda x + \kappa$). For more details see Schmickl and Crailsheim (2009); Schmickl et al. (2010).

AHHS2

Based on our experiences with AHHS1 we designed an improved variant called AHHS2. The guiding principle of this

improved controller design was to gain higher evolvability by creating smoother fitness landscapes. There were three main changes.

First, we introduced an additional rule type that implements nonlinear hormone-to-hormone interactions in the general form of $\Delta x/\Delta t = xy$, for x is the considered hormone concentration and y is the hormone concentration of the influencing hormone that triggers the considered rule. The idea is to increase the intrinsic dynamics (basically transient behavior before equilibria are reached) of the hormone network even without significant sensor input.

Second, a rule is not just triggered by exceeding or falling below a threshold but is linearly weighted within a trigger window (i.e., a tent function with a maximum of 1, defined by a center and a width, see eq. 2 below).

Third, the mutation of rule types in the form of discrete switches seemed to be too radical. This was overcome by introducing a concept of weights for rule types. Now, each rule can operate as any rule type at the same time. Each rule has a weight for each of the five rule types summing up to one (see Fig. 2). The influence of a rule type is proportional to its weight, for example, the sensor-rule aspect of a rule with a weight of 0.1 will produce only 10% of the hormone it would produce, if its weight would be 1, see $w_{\mathcal{L}}$ in eq. 1 below. A mutation will now only change two rule weights by reducing one by w and adding w to the other weight. In a well adapted controller we would expect that the weights of a rule are mainly concentrated on one or at most two rule types. Other weight distributions should be transitional only because specialization allows for better optimization.

The mathematical closed-form of this concept using the example of a linear hormone rule type is

$$\mathcal{L}(t) = w_{\mathcal{L}}\theta(H_k(t))(\lambda H_k + \kappa), \quad (1)$$

where $\mathcal{L}(t)$ is the hormone amount that is to be added to the considered hormone at time t , $w_{\mathcal{L}}$ is the weight of the linear hormone rule (see Fig. 2), k is the index of the input hormone and H_k is its concentration, λ is the dependent dose, κ the fixed dose. θ is called trigger function and defined by

$$\theta(x) = \begin{cases} \frac{1}{\eta}(\eta - |x - \zeta|) & \text{if } |x - \zeta| < \eta \\ 0 & \text{else} \end{cases}, \quad (2)$$

for trigger window center ζ and trigger window width η . For a more detailed introduction of AHHS2 and for a comparison of the AHHS approach to the standard ANN approach, see Hamann et al. (2010).

Note that the rule parameters (fixed dose, input hormone, trigger window etc.) are correlated via the rule types. For example, the input hormone is used for both the linear and the nonlinear hormone rule. If we would allow independent parameters for each rule type the genome (encoding of the controller) size would be increased by a factor of about three. This is a tradeoff in the complexity of the




Figure 2: Rule type weights of the AHHS2 approach compared to AHHS1 (abbreviations: sensor rule, linear hormone rule, nonlinear hormone rule, actuator rule).

genome and, for example, a difficulty when analyzing the results. This is related to the completeness-vs-compactness challenge (Mataric and Cliff, 1996).

Investigated scenarios

Our main focus is on the field of modular robotics and our main concern is whether we are able to evolve fast locomotion in the gait learning task. Still, we tested the AHHS approach also in an inverted pendulum task as well, due to its lower computational complexity.

Inverted pendulum

In addition to the gait learning task, we tested the AHHS approach in a task that is easier to handle: balancing the inverted pendulum (see Fig. 3). The computational demand of the gait learning task is very high due to the sophisticated simulation of physics. We satisfy the need for a simulation of lower computational complexity by introducing the inverted pendulum task. Higher statistical significance of the results can be reached within reasonable time of computation. The original inverted pendulum is only slightly related to a real robotic task. Therefore, we adapted it to our requirements. The sensors are noisy (equally distributed and uncorrelated in time, $\pm 2.3\%$) and sampling rates of sensors are low which is documented by the relation between the cycle length τ and the maximal angular velocity of $0.05\pi[1/\tau] = 9^\circ[1/\tau]$. The pendulum can move up to 9° between two calls of the controller. The controller has little time to adapt to new configurations. Furthermore, the sensors do not deliver actual angles and positions directly but partitioned onto several sensors and also relative rather than absolute (distance to wall instead of the crab's position etc.). The AHHS controls two outputs, left actuator A_0 and right actuator A_1 , while the speed control of the crab is determined by their difference. The pendulum is started in the lower equilibrium position, so the nonlinear up-swinging phase is included. Combined with the sensor noise it is impossible for the controller to balance the pendulum in the upper equilibrium position. So the task stays dynamic and the controller is exposed to new situations constantly. The

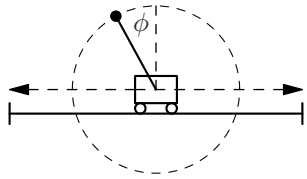


Figure 3: Inverted pendulum, pendulum free to move full 360° mounted on the crab that moves in one dimension (left/right) bounded by walls.




Figure 4: Two connected prototypes of the projects SYMBRION (2010) and REPLICATOR (2010).

fitness function is the summation over all time steps of the angular distance to the top position in radians.

Gait learning in multi-modular robotics

Gait learning in legged robotics is a commonly studied task in evolutionary robotics as reported by Nelson et al. (2009). However, here we investigate gait learning in multi-modular robotics. Each module consists of one hinge and we connect five modules. These five hinges are controlled decentrally although the modules have a low-level communication channel by means of diffusing hormones.

In contrast to the standard tasks of gait learning and collision avoidance, the challenge of gait learning in multi-modular robotics is more complex. The resulting gait is emergent due to the decentral and cooperative control of the actuators. In addition, there are several conceptionally different solutions, that is, different techniques of locomotion with good performance (e.g., caterpillar-like, erected walk, small jumps).

In each module the same controller is executed. Therefore, the gait learning task includes several sub-tasks. The organism has to break the symmetry (head and tail), synchronize through collective cooperation, and start moving into a common direction. This synchronization aspect is similar to the gait learning task for a legged robot with HyperNEAT by Clune et al. (2009).

All of this work is based on simulations as the actual hardware is not yet available (see Fig. 4 for a current prototype of Symbrion and Replicator (SYMBRION, 2010; REPLICATOR, 2010)). We use the simulation environment Symbricator3D by Winkler and Wörn (2009) that was de-

veloped for these projects. We use the current prototype design in the simulation (imported CAD data) as described in (Levi and Kernbach, 2010). However, we simplified the sensor setting to four proximity sensors (equally distributed around the robot shifted by 90 degrees: upwards, forwards, downwards, backwards). Symbricator3D is based on the game engine Delta-3D and currently uses the Open Dynamics Engine for the simulation of dynamics. The simulation of friction and momentum is important because the evolved gait behaviors rely on them. A drawback is that high computational complexity limits the number of evaluations in our evolutionary runs. We are interested in systems that evolve useful behaviors within a few hundred generations and with small populations (order of 10).

We have tested the AHHS controllers with two variants of the simulation framework. In the first version, the forces in the joints, that connect the modules, were damped and small displacements were allowed. This way the simulation reacts moderately to big forces. It turned out that caterpillar-like locomotion was favored because the damped joints support the wave motion that goes through the organism. In the second version, the joints were fully fixed. In this version of the simulation the evolution of locomotion is more difficult which will be reflected by the best fitnesses in the following.

We start the scenario with five robot modules which are simply connected in a chain. Initially this robotic organism is placed in the center of the arena. In order to increase the complexity of the gait learning task, the central area is surrounded by a low wall forming a square (its height is about half the height of a robot module). Outside the wall several cubes are placed that could only be sidestepped by the organism. An identical robot controller is uploaded to the memory of all five modules. The robot modules have to figure out their position (their role within the configuration), that is, they have to break the symmetry of the configuration in order to generate a coordinated gait. This is, for example, possible because of different outputs of proximity sensors depending on the modules' positions. There are three classes of modules defined by their characteristic sensor inputs: front module, back module, and modules in between. We use identical controllers because we want to apply them to dynamic body shapes in our future work and also a single module should have all functionality. Hence, uploading heterogeneous controllers with predefined roles would not be an option. In addition, using self-organized role assignment will allow for high scalability (using the same controller for different body sizes), plasticity (reorganization of roles in changing body shapes), and new role types might emerge that were unthought of by the human designer.

The fitness is defined by the covered distance of the organism. It is an aggregate fitness function (Nelson et al., 2009) that evaluates the organism's performance as a whole. Although the organisms might achieve advancements early in the evolutionary run, there is a bootstrapping problem. For

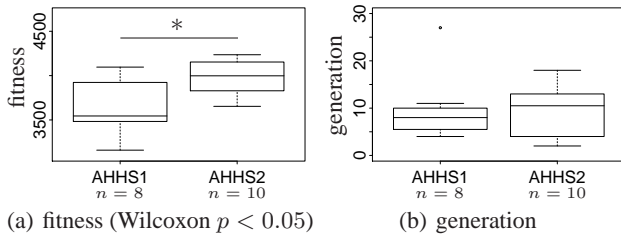


Figure 5: Inverted pendulum, AHHS1 with 60 rules, AHHS2 with 15 rules, comparison of fitness and evolution speed (generation when 75% of max. fitness was reached).

example, the downward proximity sensors will not give significant input until the organism has figured out how to erect the modules in the middle. In addition, controllers cannot evolve special techniques to climb the wall before they have actually managed to move the organism there to explore it.

Results and discussion

Inverted pendulum

The evolutionary runs of the inverted pendulum were performed with a population of 200 randomly initialized controllers. The AHHS was set to 15 hormones. For AHHS1 60 rules were used and 15 for AHHS2. The runs were stopped after 200 generations. Linear proportional selection was used and elitism was set to one. The mutation rate was 0.15 per gene with a maximal, absolute change of range 0.1. The recombination (two-point crossover) rate was 0.05.

For this task we configured AHHS with a left and a right compartment. The left compartment incorporates the left actuator A_0 , the left proximity sensor, the sensors giving the angles of the pendulum when it is in the left half etc. and for the right compartment respectively.

The comparison of the best controllers of each run is shown in Fig. 5(a). In this scenario, AHHS2 performs significantly better than AHHS1 although in terms of evolution speed there is no significant difference (see Fig. 5(b)). The AHHS2 design is the better choice in this task. The cause of the advantage of AHHS2 over AHHS1 in this task compared to the indistinct situation in the gait learning task is unclear. In future studies we will investigate whether this trend will also be observed in more complex tasks from the domain of multi-modular, evolutionary robotics.

One of the best evolved AHHS2 controllers showing interesting behavior is analyzed in the following¹. While it is not possible to keep the pendulum in the upper equilibrium for longer time due to noise, the controller still tries to maximize the time the pendulum is close to the upper equilibrium mostly by small displacements of the crab. The controller is mainly based on one hormone (H_0), and four rules (see Fig. 6). Sensor S_0 reaches its maximum, if the pendulum ap-




Figure 6: Inverted pendulum, analysis of one of the best evolved AHHS2 controllers; only most relevant rules of the evolved behavior are shown.

proaches $\phi = 0$ (top position) from the left. It triggers small displacements of the crab to the right, a behavior that keeps the pendulum turning counterclockwise with slow passes at the top position. Sensor S_9 gives the intensity of negative angular velocities of the pendulum (clockwise turns) and triggers moves of the crab to the left. The proximity sensors are not used at all. The walls are avoided by the crab movements depending on position and turning direction of the pendulum. Hence, the position of the crab is virtually encoded in the motion of the pendulum.

See Fig. 7 for the sensor, hormone, and actuator dynamics. This sample run begins with an initial ($t < 50$) move of the crab from the center to the outer left due to transient dynamics of H_0 in the left compartment (see Fig. 7(a)). This motion implements the up-swinging of the pendulum and is followed by ten small displacements of the crab to the right to keep the pendulum swinging counterclockwise. At $t = 1093$ the turning direction of the pendulum changes (see Fig. 7(b)). A sequence of right-left movements is initiated to reestablish the counterclockwise turning. Later at $t = 1933$ a phase of low angular velocity is reached which causes irregular movements of the crab that hold the pendulum close to the top position.

Gait learning

The evolutionary runs of the gait learning task were performed with a population of 20 randomly initialized controllers. The configuration of the AHHS was set to 5 hormones. The number of rules was varied between 20 and 300. The runs were stopped after 200 generations. Linear proportional selection was used and elitism was set to one. The mutation rate was 0.15 per gene (rule or hormone, with a maximal, absolute change of range 0.1). The recombination (two-point crossover) rate was 0.05. One run of the evolution (full 200 generations) took about 28 hours of CPU time (on a single core of a standard, up-to-date desktop PC). In the first version of the simulation (damped joints), the evolved behaviors reach high fitness values for all investi-

¹<http://heikohamann.de/pub/hamannEtAlAlife2010pend.mpg>




Figure 7: Inverted pendulum, most relevant hormone, sensors, and both actuator control values for both compartments (left and right) of the evolved behavior.




Figure 8: 5-module gait learning with damped joints, comparison of fitness and evolution speed, which is indicated by the generation in which 75% of the overall max. fitness ($1.41 = 0.75 \times 1.88$) was reached (if at all).

gated settings of the AHHS (see Fig. 8). Directly approaching the wall yields a fitness of about 0.7, getting one half of the modules over the wall yields a fitness of 0.8, and a fitness of above 1 is reached, if the wall is overcome. Typically the evolved behaviors rely on two or three of the five provided hormones only and make use of less than ten rules. However, a too low number of rules results in too little exploration of the behavior space. Based on preliminary tests we decided to use 30 rules for AHHS2. One AHHS2 rule is potentially active for each rule type, which corresponds to four active AHHS1 rules. However, AHHS2 cannot optimize the parameters for each rule type individually. Still, we tested the AHHS1 with 120 rules and also with a much higher number of 300 rules. The results show no statistical significant differences but show in a trend that the AHHS1 does not reach comparable results as AHHS2 with corresponding rule numbers. In addition, the behaviors evolved by AHHS1 show high variance depending on the deterministic chaos through the complex system (simulation of physics).

Using the second version of the simulation (fixed joints), we have tested smaller differences in the number of rules be-

tween AHHS1 and AHHS2. The results show that the more realistic simulation of the joints complicates the evolution of fast locomotion. However, the favoring of caterpillar-like locomotion is reduced significantly and especially in case of AHHS2 an unexpected vast diversity² of different locomotion paradigms is observed (see Fig. 9 for a short collection). Basically we observed three classes of locomotion: erected walking behavior, caterpillar-like locomotion, and locomotion through jumps. The behaviors evolved using AHHS1 were significantly less diverse. Quantifying these differences will be the focus of future studies.




Figure 9: Screenshots showing the diversity of evolved locomotion paradigms (colors represent three selected hormones in the primary colors according to the RGB color model).

The comparison of the best evolved behaviors is shown in Fig. 10(a) and the speed of evolution is shown in Fig. 10(b). 55% of the AHHS2-runs with 50 rules and 38% of the AHHS1-runs with 80 rules reach a best fitness that is within 80% of the theoretical maximum fitness of about 1.7. Significant results are only reached for AHHS1 with 20 rules compared to both AHHS1 with 80 rules and to AHHS2 with 50 rules. Noticeable is the bad performance of AHHS2 with just 20 rules both in terms of final best fitness and speed of evolution. From our observations we speculate that the initial exploration (during few of the early generations) of the

²<http://heikohamann.de/pub/hamannEtAlAlife2010.mpg>

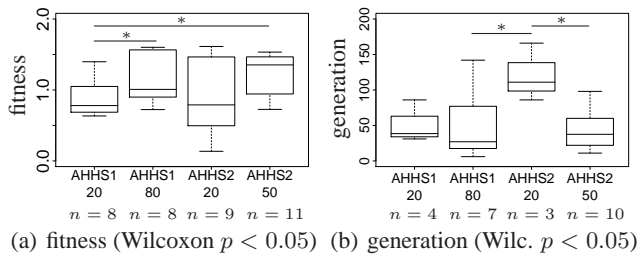


Figure 10: 5-module gait learning with fixed joints, comparison of fitness and evolution speed, which is indicated by the generation in which 75% of the overall max. fitness was reached (if at all).

search space (basically the sensory-motor configurations) is a relevant feature. Identifying the actual shortcoming of AHHS2 in this context is part of our future research.


One important aspect in the differences between the two controller types seems to be the different triggering of rules in AHHS1 and AHHS2. The behaviors of AHHS1 clearly show more fast-paced movements. With damped joints this seems to be a disadvantage as smooth movements are less likely. Using the fixed joints this sometimes results in fast locomotion through little jumps.

The evolved structures are complex and the underlying processes are often counter-intuitive. The in-depth analysis of individual behaviors is alleviated by considering the number of steps a rule has been active (triggered). Typically, about one third of the rules trigger never or very seldom.


Post-evaluation and analysis

We have investigated the behavior of one of the best evolved AHHS2 controllers in the second version of the simulator. It shows a dynamic caterpillar-like motion³. It is noticeable that the rules show characteristics of specialization and optimization. For example, often the (floating) index of the output hormone is close to an integer and often a weight for a certain rule type is above 0.5 showing the specialization of this rule. For the investigated controller we have identified three most relevant hormones: H_2 , H_3 , and H_4 . The angle of the hinge is mainly controlled by hormones H_3 and H_4 (see Fig. 11(a)). High values of H_4 turn the hinge towards $+90^\circ$ while any value of $H_3 > 0$ turns the hinge towards -90° . As a reinforcing effect there is a hormone rule that decreases H_4 , if $H_3 > 0$. H_2 shows the influence by diffusion of hormones through the organism (see Fig. 11(b)). A decreasing concentration in the back module is consequently followed by a decrease in the second last, middle, and second first module, hence, forming a hormone wave that is propagating through the organism. Finally, we investigated the influence of mutations. The leading design paradigm of AHHS2 was to improve the causality of the

³<http://heikohamann.de/pub/hamannEtAlAlife2010bestIndAna.mpg>



(a) Most relevant hormones H_3 (black) and H_4 (purple), and hinge control angle ϕ (yellow).



(b) Hormone H_2 in all five modules, demonstrating the effect of diffusion (from front module to back: light to dark).

Figure 11: 5-module gait learning with fixed joints, analysis of the evolved behavior.




Figure 12: Fitness landscape neighborhood, fitness histogram of 35 samples of mutated controllers, fitness of the original controller is for AHHS1: 0.84, for AHHS2: 0.81.

mutation operator (small changes in genome result in small changes in the behavior). This was done exemplarily by taking an evolved controller from each type. For both we produced 35 controllers by applying the mutation operator once for each. The evaluated fitnesses of these 35 controllers are shown as a histogram in Fig. 12. For AHHS1 the majority of mutated controllers had a fitness of less than 0.2. For AHHS2 the majority of mutated controllers reached about the same fitness as before. For both types some controllers reached higher fitness due variance introduced by the deterministic chaos in the simulated physics.

Conclusion and Outlook

We have reported the application of our hormone control approach to the domain of evolutionary modular robotics. The automatic synthesis of controllers, that facilitate locomotion of organisms built from five robot modules, has been effective in a majority of the evolutionary runs. Almost all evolved controllers are able to generate a form of locomotion that takes the organism at least to the wall. A majority of the evolved controllers was able to overcome the wall. An unexpected vast diversity of locomotion paradigms was evolved especially in the second version of the simulation. On the

one hand, this shows the complexity of the gait learning task in modular robotics because there are many solutions of similar utility. On the other hand, it shows the diversity of behaviors representable by the AHHS controllers.

Whether the redesigned controller AHHS2 is generally superior to the original AHHS1 design is still an open question. However, in case of the inverted pendulum it performs significantly better. In the gait learning scenario AHHS2 shows a higher diversity and behaviors with smoother movements resulting in more reliable locomotion.

There are many open issues and this research track is rather at its beginning. Our future research will include the following. The different possibilities of initializations need to be investigated extensively. For example, the controllers could be initialized with specialized sensor, hormone, and actuator rules (i.e., weights of 1). More complex tasks from the domain of modular robotics will be investigated (e.g., organisms with more modules). We plan to use environmental incremental evolution (e.g., steadily increasing heights of walls) as reported by Nakamura et al. (2000). The dynamic adaptation of rule numbers by evolution will be investigated. Hence, we will evolve hormone reaction networks through complexification similar to (Stanley and Miikkulainen, 2004). Finally, we plan to check the controllers' exploration of the sensory-motor space, especially, during the initial generations to get a better understanding of what facilitates a high diversity of solutions.

Acknowledgments

This work is supported by: EU-IST-FET project 'SYMBRION', no. 216342; EU-ICT project 'REPLICATOR', no. 216240.

References

- Chouard, T. (2010). Revenge of the hopeful monster. *Nature*, 463:864–867.
- Cliff, D., Harvey, I., and Husbands, P. (1993). Explorations in evolutionary robotics. *Adaptive Behavior*, 2.
- Clune, J., Beckmann, B. E., Ofria, C., and Pennock, R. T. (2009). Evolving coordinated quadruped gaits with the hyperneat generative encoding. In *Proceedings of the 2009 IEEE Congress on Evolutionary Computation (CEC)*. IEEE.
- Dale, K. and Husbands, P. (2010). The evolution of reaction-diffusion controllers for minimally cognitive agents. *Artificial Life*, 16(1):1–19.
- Hamann, H., Stradner, J., Schmickl, T., and Crailsheim, K. (2010). A hormone-based controller for evolutionary multi-modular robotics: From single modules to gait learning. In *Proc. of the IEEE Congr. on Evol. Computation*. accepted.
- Levi, P. and Kernbach, S., editors (2010). *Symbiotic Multi-Robot Organisms: Reliability, Adaptability, Evolution*. Springer.
- Matarić, M. J. and Cliff, D. (1996). Challenges in evolving controllers for physical robots. *Robotics and Autonomous Systems*, 19(1):67–83.
- Murata, S., Kakomura, K., and Kurokawa, H. (2008). Toward a scalable modular robotic system - navigation, docking, and integration of m-tran. *IEEE Robotics & Automation Magazine*, 14(4):56–63.
- Nakamura, H., Ishiguro, A., and Uchikawa, Y. (2000). Evolutionary construction of behavior arbitration mechanisms based on dynamically-rearranging neural networks. In *Proc. of Congr. on Evol. Computation*, volume 1, pages 158–165. IEEE.
- Neal, M. and Timmis, J. (2003). Timidity: A useful emotional mechanism for robot control? *Informatica*, 27(4):197–204.
- Nelson, A. L., Barlow, G. J., and Doitsidis, L. (2009). Fitness functions in evolutionary robotics: A survey and analysis. *Robotics and Autonomous Systems*, 57:345–370.
- Nolfi, S. and Floreano, D. (2004). *Evolutionary Robotics: The Biology, Intelligence, and Technology of Self-Organizing Machines*. MIT Press.
- REPLICATOR (2010). Project website. <http://www.replicators.eu>.
- Schmickl, T. and Crailsheim, K. (2009). Modelling a hormone-based robot controller. In *MATHMOD 2009 - 6th Vienna International Conference on Mathematical Modelling*.
- Schmickl, T., Hamann, H., Stradner, J., and Crailsheim, K. (2010). Hormone-based control for multi-modular robotics. In Levi, P. and Kernbach, S., editors, *Symbiotic Multi-Robot Organisms: Reliability, Adaptability, Evolution*. Springer.
- Shen, W.-M., Krivokon, M., Chiu, H., Everist, J., Rubenstein, M., and Venkatesh, J. (2006). Multimode locomotion via superbot reconfigurable robots. *Auton. Robots*, 20(2):165–177.
- Sims, K. (1994). Evolving 3D morphology and behavior by competition. In Brooks, R. and Maes, P., editors, *Artificial Life IV*, pages 28–39. MIT Press.
- Stanley, K. O., D'Ambrosio, D. B., and Gauci, J. (2009). A hypercube-based encoding for evolving large-scale neural networks. *Artificial Life*, 15(2):185–212.
- Stanley, K. O. and Miikkulainen, R. (2004). Competitive coevolution through evolutionary complexification. *Journal of Artificial Intelligence Research*, 21(1):63–100.
- Stradner, J., Hamann, H., Schmickl, T., and Crailsheim, K. (2009). 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)*, pages 595–600. IEEE Press.
- Stradner, J., Hamann, H., Schmickl, T., Thenius, R., and Crailsheim, K. (2010). Evolving a novel bio-inspired controller in reconfigurable robots. In *10th European Conference on Artificial Life (ECAL'09)*, LNCS. Springer. (in press).
- SYMBRION (2010). Project website. <http://www.symbion.eu>.
- Vargas, P. A., Moiola, R. C., von Zuben, F. J., and Husbands, P. (2009). Homeostasis and evolution together dealing with novelties and managing disruptions. *International Journal of Intelligent Computing and Cybernetics*, 2(3).
- Winkler, L. and Wörn, H. (2009). Symbricator3D - A distributed simulation environment for modular robots. In *ICIRA*, volume 5928 of LNCS, pages 1266–1277. Springer.