

# A Hormone-Based Controller for Evaluation-Minimal Evolution in Decentrally Controlled Systems

Heiko Hamann\*, Thomas Schmickl, and Karl Crailsheim

Artificial Life Lab

of the Department of Zoology,

Universitätsplatz 2,

Karl-Franzens University Graz,

8010 Graz, Austria

Phone: +43 316 380 3982

{heiko.hamann, thomas.schmickl, karl.crailsheim}@uni-graz.at

September 2, 2011

## Abstract

One of the main challenges in automatic controller synthesis is to develop methods that can successfully be applied for complex tasks. The difficulty is increased even more in case of settings with multiple interacting agents. We apply the Artificial Homeostatic Hormone Systems (AHHS) approach, which is inspired by the signaling network of unicellular organisms, to control a system of several independently acting agents decentrally. The approach is designed for evaluation-minimal, artificial evolution in order to be applicable to complex modular robotics scenarios. The performance of AHHS controllers is compared to NeuroEvolution of Augmenting Topologies (NEAT) in the coupled inverted pendulums benchmark. AHHS controllers are found to be better for multi-modular settings. We analyze the evolved controllers concerning the usage of sensory inputs, the emerging oscillations, and we give a nonlinear dynamics interpretation. The generalization of evolved controllers to initial conditions far from the original conditions is investigated and found to be good. Similarly the performance of controllers scales well even with module numbers different from the original domain the controller was evolved for. Two reference implementations of a similar controller approach are reported and shown to have shortcomings. We discuss the related work and conclude by summarizing the main contributions of our work.

**Keywords:** controller synthesis, hormone/endocrine system, decentral control, evolutionary robotics, inspiration by signaling networks of unicellular organisms

---

\*corresponding author

# Contents

<b>1</b>	<b>Introduction</b>	<b>3</b>
1.1	Challenges of controller synthesis . . . . .	4
1.2	Inspiration by signaling networks in unicellular organisms . . . . .	5
<b>2</b>	<b>Artificial Homeostatic Hormone Systems (AHHS)</b>	<b>7</b>
2.1	Mathematical model of AHHS . . . . .	7
2.2	Encoding AHHS in a genome . . . . .	11
2.3	Discussion of several design decisions . . . . .	14
2.4	Reduction of computational costs . . . . .	17
<b>3</b>	<b>Domain: Coupled inverted pendulums</b>	<b>18</b>
3.1	Description of the coupled inverted pendulums benchmark . . . . .	18
3.2	Discussion of properties of the coupled pendulums benchmark . . . . .	19
<b>4</b>	<b>Results</b>	<b>21</b>
4.1	Comparison of AHHS to NEAT . . . . .	22
4.2	Applying noise during evolution for better performance . . . . .	23
4.3	Analysis of an evolved controller . . . . .	24
<b>5</b>	<b>AHHS – Mode of operation</b>	<b>25</b>
5.1	Influence of rule numbers . . . . .	28
5.2	Most influential sensors . . . . .	28
5.3	Nonlinear dynamics interpretation . . . . .	29
5.4	Control based on oscillations . . . . .	30
<b>6</b>	<b>Generalization and Scaling</b>	<b>32</b>
6.1	Generalization to initial conditions far from original domain . . . . .	32
6.2	Scaling of module numbers . . . . .	33
<b>7</b>	<b>Comparison to similar implementations</b>	<b>34</b>
7.1	Direct encoding of lookup tables . . . . .	35
7.2	Lookup tables based on Fourier series . . . . .	36
<b>8</b>	<b>Related work</b>	<b>36</b>
8.1	Artificial neural networks, NEAT, HyperNEAT, and GasNet . . . . .	36
8.2	Gene regulatory networks and reaction-diffusion systems . . . . .	38
<b>9</b>	<b>Conclusion</b>	<b>39</b>

# 1 Introduction

The survival of natural creatures and the success of artificial creatures is importantly affected by their cognitive abilities. An embodied, mobile, autonomous agent has to process sensory inputs and has to control its actuators appropriately in order to survive or to fulfill its assigned task. The number of possible actions, which could potentially be performed by such an agent, is generally high. Also the number of constraints, that confine the feasibility of actions, is high due to limited capabilities of the agents' bodies or induced by the environment.

The organs in natural agents, that select appropriate actions based on perceptions, emerged through processes of natural selection. The variety of naturally evolved control systems is vast and goes much beyond the frequently stated example of the central nervous system of vertebrates. This is most evident in case of unicellular organisms that often show non-trivial behavior without possessing a single nerve cell, for example, *Paramecium* [2, 6].

Generally, nature has evolved several principles of communication that partially act in parallel within organisms. In unicellular organisms (microorganisms, microbes), receptors can alter the production of chemical cell signals, which diffuse within the cell and integrate. These signals do not only transport information through space because they are sometimes part of complex biochemical cascades. Some of these signals also allow the organism to switch between different modes of operation (e.g., motion principles or physiological states) as for example in *Paramecium* [6]. As cells are internally structured by compartmentalization, these processes can be interpreted as biochemical computation in space and time [4, 5, 6, 48]. Generally this sort of communication can be interpreted as broadcast communication because there is no direct link between sender and receiver, nor is there a dedicated addressed message. However, as intracellular chemical signals usually affect specific organelles/actuators and as specific compositions of chemical signals are often correlated with specific cell states, an emitter-channel-receiver model can still be applied to this sort of communication.

Within cells of multi-cellular organisms (metazoa), several pathways of communication exist, which differ concerning the processes they exploit. Between neighboring cells trans-membrane proteins act as important mechanisms to transport chemical molecules against the concentration gradient, which makes this sort of communication significantly different from basal diffusion. The most prominent examples are, for example, ion pumps.

Generally, cell-to-cell signaling can be divided into three categories [1]: Communication between cells based on direct contact and strong dependence on local morphology (*juxtacrine signaling*), communication over short distances and medium dependence on local morphology (*paracrine signaling*), and communication over large distances and/or scales and weak dependence on local morphology (*endocrine signaling*). Gap junctions and notch signaling, belonging to the first category, bind communication pathways strongly to the local morphology around the sender's location. Paracrine hormones and neuro-transmitters, still having a high linkage between morphology (location of the communicat-

ing cells) and communication functionality, belong to the second category. In the third category, endocrine hormones allow communication over greater distances and show only a weak coupling between morphology and communication functionality.

Returning to the main focus of this paper, we interpret juxtacrine signaling and paracrine signaling as belonging to one group of communication mechanisms, as there exists a clear morphological (topology-mediated) coupling between sender and receiver. Also far-reaching neural communication can be classified belonging to this group, as neural communication follows a dedicated pathway within the nervous system. We argue that the concept of Artificial Neural Networks (ANN) is mainly inspired by this group of directed communication (unicast) which is also evidently indicated by the directed edges between neurons.

In addition, we interpret intracellular signal processing and endocrine communication as belonging to a second group of communication mechanisms, as these communication principles exhibit a rather loose morphological coupling (except receptor-mediated linkage) between sender and receiver. These processes are build on diffusion of chemical signals. The design principles of the Artificial Homeostatic Hormone System (AHHS) clearly reflect this group of natural broadcast communication (undirected communication).

It is to be noted that presenting these two groups as distinct paradigms is a simplification because they are two extremes in a continuous transition between paradigms of biological communication. For example neuro-endocrine cells and trans-membrane proteins are features that add aspects of one group to the mechanisms of the other group. Also other means of communication or communication clues such as mechanical forces, temperatures, or external light inputs might be important in inter- and intracellular communication.

Reaching out for new sources of inspiration in the context of engineering controllers of modular technical devices is worthwhile, although the high capabilities especially of the human brain are unquestioned. However, the artificial synthesis of controllers of that complexity are out of reach. Actually, the synthesis of any controller that sufficiently completes the information processes for an agent living in a dynamic, unpredictable environment is challenging and the non-neuronal information processing of unicellular organisms definitely generates non-trivial behaviors [4].

## 1.1 Challenges of controller synthesis

The complexity of controller synthesis can be rated, for example, by investigating the field of evolutionary robotics [44, 19, 69]. The (semi-)automatic synthesis of robot controllers by applying artificial evolution belongs to the software section of evolutionary robotics [12]. The curse of complexity bears the main challenge in this field. An increase in the difficulty of the desired behavior seems to result in a significant increase in the complexity of its evolution. This is partially documented by the absence of complex benchmark tasks in the literature [42]. It is still an open question of how to find either a basic principle of the controller design that is appropriate and universal or one that is optimal for

a given scenario [74, 21, 22, 7, 8]. The latter seems to be more likely to succeed because the existence of a true general problem solver is very questionable (cf. the ‘no free lunch’ theorem [76]). This is, however, intuitively in conflict with the diversity of the human problem solving capabilities indicating that a kind of non-exhaustive ‘almost general’ problem solver exists.

Translating the concept of controller design and general problem solving to the context of evolutionary algorithms [29, 34, 50, 58] is associated with the problem of generating general or scenario-dependent, smooth fitness landscapes. The general controller design defines 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 density of acceptable solutions in search space should be kept high, the fitness landscape should generally be smooth with a minimum number of local optima. Experience shows that these two criteria are contradicting. In addition, the search space is usually high-dimensional and might have unfavorable structures. We summarize this set of challenges by the aim to ‘strive for high evolvability’.

One focus of our research track is to design fitness landscapes by applying appropriate controller designs. 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 [24]. However, whether high causality is really desirable, is questionable (e.g., cf. Chouard [11]).

The challenge of appropriate, efficient, and safe behavioral control is brought to a new level whenever groups of agents interact [78]. This holds also for organisms that are built from autonomous entities which collectively establish (possibly several different) topologies of their bodies. Examples of natural systems that belong to this category are the slime moulds *Dictyostelium discoideum* and *Dictyostelium mucoroides* [9] or the algae *Volvox aureus* [30]. *Volvox aureus* is living in spherical aggregates consisting of hundreds or thousands of cells. The colony is able to perform collective phototaxis to stay in favorable light conditions. This is achieved by modulating the flagellar activity of cells depending on the local light intensity [26]. The responses to the light stimuli vary around the sphere forming a gradient of flagellar activity suggesting inter-cellular interaction (be it chemical, electrical or mechanical). This allows global coordination throughout the colony. Such organisms and behaviors are good sources of inspiration for paradigms of decentral control.

## 1.2 Inspiration by signaling networks in unicellular organisms

Especially the utilization of inter-cellular interactions are of interest here. Organism such as *Volvox aureus* have not evolved explicit communication processes between a sender and a receiver (unicast) but rather implement implicit communication by chemical gradients (broadcast) or mechanical clues. Based on these considerations and inspired by the signaling network in unicellular organisms we have proposed controllers based on Artificial Homeostatic Hormone Systems (AHHS) before [24, 23, 55, 56, 54, 57, 66, 65]. The wording ‘hormone’ is meant

as a generic term for any kind of cell signals. These systems can be viewed as reaction-diffusion systems that are embedded into the autonomous agents. Sensory stimuli are converted into hormone secretions (cell signals) that, in turn, control the actuators. In addition, hormones interact linearly and nonlinearly comparable to the hidden layer of ANN. Such systems show homeostatic processes because they typically converge to trivial steady states for constant sensory input. Sensory stimuli trigger hormone secretion, hormone concentrations are basically integrated (a form of memory), and decomposed over time (oblivion). However, during a limited period of time (transient) after a stimulus they also show variant behavior, especially, if nonlinear hormone-to-hormone interactions are applied. This way the bootstrapping problem of how to generate many sensory-motor configurations initially is overcome because the controllers explore many such configurations even without input. The concept of AHHS is related to gene regulatory networks. However, here each edge has its own activation threshold and redundant edges with different activations are allowed. Hormones may diffuse within virtual, internal, spatial structures in the agent but also from one agent to neighboring agents.

The desired main application of AHHS is multi-modular robotics [67, 51]. In this field, autonomous robotic modules are studied, that are able to physically connect to each other, and they can also establish a communication and energy connection. Hence, they form a super-robot ('organism'), that is able to re-configure its body shape, see for example [61, 53, 59, 39]. 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 [44]. Although we evolve cooperative behaviors by evolving a kind of self-organized role selection, there is no co-evolution due to identical genomes in every module.

In the following we give a detailed description of AHHS, discuss several design decisions and how the computational complexity can be limited. The benchmark based on inverted pendulums, that is applied throughout the paper, is described and discussed. In the second half of the paper, we investigate many properties of AHHS, such as the resulting control networks, how they compare to controllers evolved by applying NEAT [63], the nonlinear dynamics, the oscillations in such systems, the generalization of evolved controllers to initial conditions far from the original domain and their scaling concerning increased and decreased complexity of the scenario. Finally we report two reference implementations, discuss their shortcomings, and discuss related work.

## 2 Artificial Homeostatic Hormone Systems (AHHS)

The original version of the artificial homeostatic hormone systems (AHHS) was introduced by Schmickl and Crailsheim [54], Schmickl et al. [57]. Several publications followed that report applications of AHHS to evolution of controllers for simulated single robots [66] and hand-coded controllers for real single robots [65, 56]. General perspectives of the AHHS approach were reported by Schmickl et al. [55]. The authors reported an improved version of the AHHS in [23], which will be used in this work (the old version is not used here). An application to modular robotics is reported in [24].

The basic concept of AHHS is described in the following (for a less detailed description see [23]). An AHHS is defined by a set of hormones, that define general properties of hormones, and a set of rules that define how sensory input and hormone concentrations are converted into hormone concentration changes and actuator control signals. The hormones have base production rates, that is, they are allowed to increase their concentrations independently. A decay rate defines their independent decrease of concentration. In addition, rules manipulate these hormones and also influence the change of hormone concentrations based on sensory input (sensor sub-rules) and based on the concentrations of other hormones or the hormone itself (hormone-to-hormone sub-rules). Finally, the actuators are controlled by rules that transfer hormone concentrations into actuator control values.

The diffusion of hormones is defined based on any spatial/modular topology. For example, the topology could be an abstract compartment structure within the agent in allusion to compartmentalizations in biological cells [68] or a loosely coupled aggregate of several agents. Hence, diffusion generates a means of communication which is however rather coarse. It implements a communication based on broadcasts, which could also be called ‘implicit communication’ in order to distinguish it from direct communication (unicast or explicit communication) of a preserved message between sender and receiver.

### 2.1 Mathematical model of AHHS

Now we give a detailed description of AHHS. For each hormone its special characteristics are defined, such as decay rate, base production rate, and diffusion coefficient (see Table 1 for a summary). A hormone  $H$  is defined by a 3-tuple

$$G^H = (\alpha, \mu, D), \quad (1)$$

for base production rate  $\alpha$ , decay rate  $\mu$ , and diffusion coefficient  $D$ . This tuple is also called ‘hormone gene’ below because we apply genetic algorithms to evolve AHHS.

The basic concept of rules is that each rule contains four sub-rules, that is, each rule defines a package of sub-rules: actuator sub-rule, sensor sub-rule, linear hormone-to-hormone sub-rule, and nonlinear hormone-to-hormone sub-rule. Nonlinear sub-rules were introduced in addition to linear interactions ‘to




Figure 1: Sensor-to-hormone, hormone-to-hormone, and hormone-to-actuator interactions defined by a single rule. The two arrows between actuator-hormone pair and sensor-hormone pair indicate the two-fold interaction by the linear and the nonlinear sub-rule.

allow intrinsic dynamics of higher complexity' [23]. Each rule is listening for one selected sensor and influences one selected actuator. However, a direct influence from sensor to actuator is generally excluded because the hormones, that are influenced by the sensor, do not influence the actuator themselves (see Fig. 1). Each rule influences a pair of hormones through sensory input, it influences an actuator based on the concentration of a pair of hormones, and it influences one pair of hormones based on the concentrations of another (or possibly the same) pair of hormones. The pairs of hormones are defined by weighted (floating) indices, that is, a real number. Floating indices were introduced to optimize the local search when synthesizing AHHS with evolutionary methods as explained in Sec. 2.3.4. For example, a floating index of 0.5 addresses with equal weights of 50% hormone  $H_0$  and hormone  $H_1$ . A floating index of 1.9 would address  $H_2$  with a weight of 90% and  $H_1$  with 10%. Hence, a rule defines a package of sub-rules that can be viewed as a (sub-)graph or network as sketched in Fig. 1. A rule is defined by a 12-tuple (also called 'rule gene' below)

$$G^R = (w^A, w^S, w^L, w^N, \zeta, \eta, \lambda, \kappa, s, a, h, k), \quad (2)$$

for sub-rule weights  $w$ , trigger window center  $\zeta$ , trigger window width  $\eta$ , dependent dose  $\lambda$ , fixed dose  $\kappa$ , sensor ID  $s$ , actuator ID  $a$ , input hormone ID  $h$  defining the sensor-hormone pair, and output hormone ID  $k$  defining the actuator-hormone pair. Each of the sub-rules has a weight determining its influence. There is also an implicit weight  $w^I$  defining a virtual 'idle sub-rule', that corresponds to a sub-rule without effect. All sub-rule weights sum up to  $\sum_{i \in \{A, S, L, N, I\}} w^i = 1$ . Hence, an increasing influence of one sub-rule always comes with a decreasing influence of another sub-rule. It is possible to deactivate a rule by setting  $w^I = 1$  and it is also possible to define a specialized rule by setting one of the other weights to one. Such a rule would be fully dedicated to one of the sub-rules (e.g., a sensor rule with  $w^S = 1$ ). This concept of weights allows a step-by-step (in principle continuously) transition of rules between the rule types which is important for high evolvability. Note that the parameters of



the sub-rules cannot be separated in individual sub-rule tuples because the sub-rules share several parameters (detailed motivation and discussion in Sec. 2.3). In the following we define the hormone dynamics of AHHS. The dynamics of hormone concentrations is influenced by a constant production of hormone, diffusion of hormones, decay, and the summation of influences by other hormones (and itself) and sensors as defined by rules. Note that diffusion of hormone concentrations within a certain spatial structure of compartments (see Sec. 2.3.3) establishes the only means of communication in this system which is discussed in Sec. 2.3.2. The change of hormone concentration  $H_h^c$  of hormone  $h$  in compartment  $c$  at time  $t$  is described in dimensionless units by

$$\begin{aligned} \frac{\Delta H_h^c}{\Delta t} = & \alpha_h + D_h \nabla^2 H_h^c(t) - \mu_h H_h^c(t) \\ & + \sum_i \mathcal{S}_i(t) + \mathcal{L}_i(t) + \mathcal{N}_i(t), \end{aligned} \quad (3)$$

for the parameters of hormone  $H_h$  as defined in the hormone gene (eq. 1): production rate  $\alpha_h$ , diffusion coefficient  $D_h$ , decay rate  $\mu_h$ , and the summed influence of all applicable sensor sub-rules  $\mathcal{S}_i$ , applicable linear hormone sub-rules  $\mathcal{L}_i$ , and applicable nonlinear hormone sub-rules  $\mathcal{N}_i$  (all of which are introduced in the following). The rationale for defining the influence of sensors and hormones by sums of several ‘sub-influences’ defined by sub-rules is motivated and discussed in Sec. 2.3.1. The diffusion term is noted continuously to obtain a concise description although it is, of course, discretized in the implementation. Hormones have common minimal ( $H_{\min}$ ) and maximal values ( $H_{\max}$ ) that cannot be exceeded (i.e.,  $\forall h, t : H_{\min} \leq H_h(t) \leq H_{\max}$ ) to have well defined intervals for the dynamics of hormones.


We define the sensor sub-rule, which defines the influence to the concentration  $H_h^c$  of a hormone  $h \in \mathbb{R}^+$  (i.e., each sensor sub-rule is applied to two output hormones  $h_1$  and  $h_2$  because we use weighted IDs of hormones and get hormone pairs, cf. Fig. 1) in compartment  $c$  (cf. eq. 3),

$$\mathcal{S}_i(t) = w_i^S \theta(S_s(t)) (\sigma_s S_s(t) \lambda_i + \kappa_i), \quad (4)$$


for sensor sub-rule weight  $w_i^S$ , sensory input  $S_s(t)$  from sensor  $s$  (subscript specified by rule gene, cf. eq. 2), a linear sensor scaling constant  $\sigma_s$ , dependent dose  $\lambda_i$ , and fixed dose  $\kappa_i$  (described in the following). We use a tent function  $\theta$  as a trigger function that determines whether and with which intensity the sub-rule is executed. This was introduced in favor of a mere threshold because experiments indicated better performance [24]. It is defined by

$$\theta(x) = \begin{cases} \frac{1}{\eta_i}(\eta_i - |x - \zeta_i|) & \text{if } |x - \zeta_i| < \eta_i \\ 0 & \text{else} \end{cases}, \quad (5)$$


for trigger window center  $\zeta_i$  and trigger window width  $\eta_i$ . This is a linear weighting depending on the distance to the trigger window center and its width. Examples in Fig. 2 show the interaction of eqs. 4 and 5.




(a)  $\zeta = 0.4$ ,  $\eta = 0.2$  and settings of  $\lambda, \kappa$  as in the corresponding figures.



(b) Based on slope  $\lambda = 1$  and y-intercept  $\kappa = -0.2$  of straight line  $b'$  shown in Fig. a.



(c) Based on slope  $\lambda = 1$  and y-intercept  $\kappa = -0.4$  of straight line  $c'$  shown in Fig. a.



(d) Based on slope  $\lambda = -1$  and y-intercept  $\kappa = 0.5$  of straight line  $d'$  shown in Fig. a.

Figure 2: Examples of parameter settings for  $\zeta, \eta, \lambda, \kappa$  defining trigger function  $\theta$  and contribution  $\mathcal{S}$  of a sensor sub-rule ( $\sigma_s = 1, w^S = 1$ ). Figures b, c, and d are the result of multiplying trigger function  $\theta$  as shown in Fig. a with each of the straight lines  $a', b',$  and  $c'$  in Fig. a (cf. eqs. 4 and 5).

We define the linear hormone sub-rule

$$\mathcal{L}_i(t) = w_i^{\mathcal{L}} \theta(H_k(t))(H_k(t)\lambda_i + \kappa_i), \quad (6)$$

which is applied to an output hormone concentration  $H_h^c$  (to each of the sensor hormone pair as shown in Fig. 1) in each compartment  $c$  (cf. eq. 3). The input is hormone concentration  $H_k$ . Note that  $h = k$  is allowed, thus self-referencing of a hormone is possible. This is an important feature because it introduces a feedback that can be leveraged to generate complex dynamics (e.g., see Sec. 5.3). The other parameters are as defined above.

The nonlinear hormone sub-rule is characterized by the product of the input and the output hormone concentrations  $H_k(t)H_h(t)$  (this choice is remotely inspired by chemical reactions and their rate equations) and defined by

$$\mathcal{N}_i(t) = w_i^{\mathcal{N}} \theta(H_k(t))(H_k(t)H_h(t)\lambda_i + \kappa_i), \quad (7)$$

which is applied to an output hormone  $H_h^c$  in each compartment  $c$ .

We define the actuator sub-rule

$$\mathcal{A}_i(t) = w_i^{\mathcal{A}} \theta(H_k(t))(H_k(t)\lambda_i + \kappa_i), \quad (8)$$

which is applied to an actuator control  $A_a$  (cf. eq. 9). The current actuator control value  $A_a$  of actuator  $a$  in time step  $t$  is defined by

$$A_a(t) = \begin{cases} A_a^{\min}, & \text{if } \sigma_a \sum_i \mathcal{A}_i(t) \leq A_a^{\min} \\ \sigma_a \sum_i \mathcal{A}_i(t), & \text{if } A_a^{\min} < \sigma_a \sum_i \mathcal{A}_i(t) < A_a^{\max} \\ A_a^{\max}, & \text{if } \sigma_a \sum_i \mathcal{A}_i(t) \geq A_a^{\max} \end{cases}, \quad (9)$$

for actuator sub-rules  $\mathcal{A}_i$ , maximum actuator value  $A_a^{\max}$ , minimum actuator value  $A_a^{\min}$ , and actuator scaling constant  $\sigma_a$  that linearly scales hormone values to the relevant actuator control value interval. Eq. 9 implements the control of the actuator by the summed influence of all actuator sub-rules and the limitation to the relevant interval.


## 2.2 Encoding AHHS in a genome

A specific data structure, which parametrizes the AHHS controllers, is introduced that is used by the genetic algorithm as genome. The genome is a 2-tuple  $\Gamma = (C_h, C_r)$  consisting of two logical entities: *hormone chromosome*  $C_h$  and *rule chromosome*  $C_r$ . There is one hormone gene  $G^H$  for each of the  $N$  hormones in the hormone chromosome  $C_h = (G_1^H, G_2^H, \dots, G_N^H)$ . There is one rule gene  $G^R$  for each of the  $M$  rules in the rule chromosome  $C_r = (G_1^R, G_2^R, \dots, G_M^R)$ . In Table 1 a listing of all genes in these two types of chromosomes is given. The hormone genes contain the actual parameters in a 3-tuple (see eq. 1). The rule genes contain the actual parameters in a 12-tuple (see eq. 2). Hence, the genome can be implemented, for example, as a mere array of floating point numbers. A short summary of the AHHS is shown in Fig. 3.

Table 1: The genome of the AHHS controller.

<b>Hormone Chromosome</b>		
Gene	Description	Range
base production rate $\alpha$	amount that is produced without sensory stimulation	$0 \leq \alpha \leq H_{\max}$
decay rate $\mu$	cf. eq. 3	$0 \leq \mu \leq 1$
diffusion coefficient $D$	cf. eq. 3	$0 \leq D \leq 1$
max./min. value of hormone $H_{\max}, H_{\min}$	value at which a saturation/bottom is forced	$H_{\max}, H_{\min} \in \mathbb{R}$
<b>Rule Chromosome</b>		
Gene	Description	Range
sub-rule type weight $w$	weights of the three sub-rules and the idle sub-rule $w^I$	$w^S + w^A + w^H = 1 - w^I, 0 \leq w \leq 1$
trigger window center $\zeta$	defines (along with trigger window width $\eta$ ) the triggering condition and intensity	$H_{\min} \leq \zeta \leq H_{\max}$
trigger window width $\eta$	defines (along with trigger window center $\zeta$ ) the triggering condition and intensity	$H_{\min} \leq \eta \leq H_{\max}$
dependent dose $\lambda$	cf. eqs. 8, 4-7	$\lambda \in \mathbb{R}$
fixed dose $\kappa$	cf. eqs. 8, 4-7	$\kappa \in \mathbb{R}$
sensory input $s$	weighted IDs of the sensors that influence the hormone through sensor sub-rules $\mathcal{S}$	$\mathbb{R}^+$ (limited by sensor number)
actuator output $a$	weighted IDs of the actuators that are influenced by the hormone through actuator sub-rules $\mathcal{A}$	$\mathbb{R}^+$ (limited by actuator number)
hormone input $h$	weighted IDs of the hormones that are influenced through sensor sub-rule $\mathcal{S}$ or that influences another hormone through hormone sub-rules $\mathcal{H}$	$\mathbb{R}^+$ (limited by hormone number)
hormone output $k$	weighted IDs of the influenced hormones (through hormone sub-rules $\mathcal{H}$ )	$\mathbb{R}^+$ (limited by hormone number)

## Artificial Homeostatic Hormone Systems (AHHS) in a nutshell



An AHHS genome is a 2-tuple  $\Gamma = (C_h, C_r)$  consisting of a hormone chromosome  $C_h = (G_1^H, G_2^H, \dots, G_N^H)$  and a rule chromosome  $C_r = (G_1^R, G_2^R, \dots, G_M^R)$ . A hormone is a 3-tuple  $G^H = (\alpha, \mu, D)$  and a rule is a 12-tuple  $G^R = (w^A, w^S, w^L, w^N, \zeta, \eta, \lambda, \kappa, s, a, h, k)$ .

Change of hormone concentration:  

$$\frac{\Delta H_h^c}{\Delta t} = \alpha_h + D_h \nabla^2 H_h^c(t) - \mu_h H_h^c(t) + \sum_i S_i(t) + \mathcal{L}_i(t) + \mathcal{N}_i(t)$$

Sensor sub-rule:

$$S_i(t) = w_i^S \theta(S_s(t)) (\sigma_s S_s(t) \lambda_i + \kappa_i)$$

Trigger function:

$$\theta(x) = \begin{cases} \frac{1}{\eta_i} (\eta_i - |x - \zeta_i|) & \text{if } |x - \zeta_i| < \eta_i \\ 0 & \text{else} \end{cases}$$

Linear hormone sub-rule:

$$\mathcal{L}_i(t) = w_i^L \theta(H_k(t)) (H_k(t) \lambda_i + \kappa_i)$$

Nonlinear hormone sub-rule:

$$\mathcal{N}_i(t) = w_i^N \theta(H_k(t)) (H_k(t) H_h(t) \lambda_i + \kappa_i)$$


Actuator sub-rule:

$$\mathcal{A}_i(t) = w_i^A \theta(H_k(t)) (H_k(t) \lambda_i + \kappa_i)$$


$\alpha$	base production rate
$\mu$	decay rate
$D$	diffusion coefficient
$H_h^c$	hormone concentr. of h. $h$ in compart. $c$
$w$	sub-rule weight
$\sigma_s$	sensor scaling constant
$\sigma_a$	actuator scaling const.
$\lambda$	dependent dose
$\kappa$	fixed dose
$\eta$	trigger window width
$\zeta$	trigger window center
$s$	sensor ID
$a$	actuator ID
$h$	input hormone ID
$k$	output hormone ID

Actuator control value:

$$A_a(t) = \begin{cases} A_a^{\min}, & \text{if } \sigma_a \sum_i \mathcal{A}_i(t) \leq A_a^{\min} \\ \sigma_a \sum_i \mathcal{A}_i(t), & \text{if } A_a^{\min} < \sigma_a \sum_i \mathcal{A}_i(t) < A_a^{\max} \\ A_a^{\max}, & \text{if } \sigma_a \sum_i \mathcal{A}_i(t) \geq A_a^{\max} \end{cases}$$



(a) network of a rule package



(b) trigger function, (c) contribution of the line  $\lambda = -1$ ,  $\kappa = 0.5$  sub-rule

Figure 3: Short summary of AHHS.

## 2.3 Discussion of several design decisions

The proposed design of AHHS controllers includes several fundamental decisions that are discussed in the following.

### 2.3.1 Indirect shaping of sensor-to-hormone, hormone-to-hormone, and hormone-to-actuator mappings

One of the main features of the AHHS approach is the definition of behaviors through indirect shaping of mappings  $f_s$  of sensor values to hormone concentration changes (sensor-to-hormone mappings, see eq. 3 and Fig. 4)

$$f_s : (S_0, S_1, \dots) \mapsto \Delta H, \quad (10)$$

$$f_s(S_0, S_1, \dots) = \sum_i \mathcal{S}_i(t), \quad (11)$$


based on the summation of several sensor-to-hormone functions  $\mathcal{S}_i(t)$  (see eq. 4 and Fig. 2) and similarly for hormone concentrations to hormone concentrations (hormone-to-hormone mappings) and hormone concentrations to actuator values (hormone-to-actuator mappings). Only the summands  $\mathcal{S}_i(t)$  can be directly changed, for example, by methods of evolutionary computation. This seems to be a performance-sensitive feature of AHHS as shown below in Sec. 7 by the comparison to reference implementations.

Defining the controllers via summations of several functions (cf. eq. 3 and Fig. 4) rather than using one function description alone has proved to be advantageous as shown below by empirical evidence. The number of functions (i.e., the number of rules) that contribute, for example, to a hormone-to-hormone mapping varies depending on how many rules apply to the considered hormone concentration pair  $H_h$  and  $H_k$  (similar for sensor-to-hormone and hormone-to-actuator sub-rules). Fig. 4(a) shows such a situation where 7 rules apply. The resulting hormone-to-hormone mapping defines the hormone concentration change  $\Delta H_{\text{out}}$  based on the hormone concentration  $H_{\text{in}}$  of an input hormone. Note that small changes of a single function will also have small effects on the summation. Especially, qualitative, global changes of the summation are unlikely. If this is really desirable is questionable as mentioned above [11].


If we investigate the summation of linear and nonlinear hormone-to-hormone sub-rules we get 2-dimensional functions for the case of  $h \neq k$  (i.e., input and output hormone index are different) as shown in Fig. 4(b).

The summed influences of sub-rules can be interpreted as those entities that define the resulting behaviors and on which the artificial evolution acts. This is, for example, shown in Fig. 4(c). In a process over many generations the general quality of the hormone-to-hormone mapping is preserved but the quantities are changed (in this case almost linearly) over generations.


A comparison to ANN is self-evident because AHHS controllers can be interpreted as hormone reaction networks (cf. Fig. 1 and Fig. 9). The equivalent



(a) Typical linear hormone-to-hormone mapping (dashed) based on summation of 7 sub-rules (solid).



(b) Hormone-to-hormone mapping including both linear and nonlinear sub-rules. The plot is 2-dimensional because the nonlinear hormone-to-hormone mapping has both  $H_{in}$  and  $H_{out}$  as input according to eq. 7:  $\mathcal{N}_i(t) = w_i^N \theta(H_{in}(t))(H_{in} H_{out} \lambda_i + \kappa_i)$ . The different shades of gray display the change of hormone concentration for hormone  $H_{out}$ .



(c) Evolution of one hormone-to-hormone mapping, numbers give number of generation (see Table 3 for used parameters).

Figure 4: Summations of the influences by hormone-to-hormone sub-rules define overall hormone-to-hormone mappings (similar for sensor-to-hormone and hormone-to-actuator sub-rules).

to the above mentioned summed X-to-hormone mappings in ANN is not well-defined but could be a (rather exceptional) network function of the form

$$\text{net}_j(t) = \sum_i w_{ij}(o_i(t)), \quad (12)$$

with  $w_{ij}$  are weight *functions* depending on the outputs  $o_i$  of connected neurons instead of mere weight constants. These weight functions would be subject to evolution. The authors are not aware of any approach pursuing this method with methods of evolutionary computation. Still, this approach is representable with standard ANN by implementing the functions  $w_{ij}$  through according sub-trees (or sub-graphs) of neurons with regular weights. It is, however, questionable how such structures should emerge in artificial evolution—a process that is principally blind to topological qualities of the network. Hence, one could argue that the AHHS approach predefines useful clusters of features that would correspond to independent sub-graphs in ANN.

Another difference between AHHS and ANN is the analogy of a ‘conservation of mass’ in case of AHHS. Hormone concentrations are produced, preserved over time, and vanish only obeying a regular decay process. Hence, the emergence of processes based on memorized sensory inputs is likely. In contrast, the activations of neurons in ANN are not explicitly preserved which can only be implemented by establishing feedback loops.

### 2.3.2 Fuzzification and low-level communication

The internal state of an AHHS controller is determined by a few continuous values (one for each hormone). The agents’ behaviors are determined by these values in combination with the hormone-to-actuator mappings. For small deviations these mappings give typically similar results (see Fig. 9) and introduce a kind of fuzzification into the system. Hence, diffusion of hormone concentrations implements a low-level communication (implicit communication) because hormone values are propagated through the compartment/module system and similar hormone values have similar ‘meanings’ within the controller. This might prove to be a powerful process because, on the one hand, communication is important in decentral control and, on the other hand, the alternative of evolving sender-receiver pairs would be a difficult problem [16, 28, 17].

### 2.3.3 Spatial structure due to compartments

The idea of compartmentalizing either a single agent into abstract compartments or just using the natural compartmentalization (e.g., in modular robotics) is to allow a reflection of the agents’ embodiment in the controller. This way sensors, actuators, and possibly even sub-rules (e.g., sensor sub-rules via their coupling with certain sensors) are associated with certain compartments and allow a body-based (embodied) modularization of the controller. In addition, the evolution of complex compartmentalizations consisting of dozens of compartments might enable behaviors that are otherwise difficult to evolve. This is, however, part of our future work agenda.



### 2.3.4 Floating indices and maximal causality of mutations

The idea of floating indices to address hormones and the idea of using weights to specify the influence of sub-rules is to maximize the causality of the mutation operator by reducing the impact of a single mutation to the agent’s resulting behavior. Using discrete indices would correspond to implement mutations as switches. Changing a rule from one hormone to another or changing from one sub-rule to another corresponds typically, however, to a radical change in the behavior. Radical changes are typically fatal in the evolutionary process. Having many fatal mutations complicates the search and the genetic algorithm might be caught in local maxima more often. However, radical changes might be a relevant process of natural evolution [11]. Hence, a good choice should depend on the fitness landscape of the investigated domain. However, research of rugged fitness landscapes [45] that investigate fitness landscapes, that are more complex than Kauffman’s NK model [33], is still pending mainly due to the computational complexity of empirical studies.

### 2.3.5 Dependent parameters in rules

The idea of packaging several sub-rule types in one rule is a tradeoff. Several parameters of a rule, such as the dependent dose  $\lambda$  or the fixed dose  $\kappa$ , are used for several different sub-rule types, that is, these values are correlated and cannot be optimized independently by the applied method of evolutionary computation unless the corresponding weights are decreased close to zero. The independent parameters (e.g., sensor ID or actuator ID), in turn, can become silent genes, if the corresponding sub-rule weights are zero. An alternative would be a considerable increase of the genome size (e.g., a full set of parameters for each sub-rule). However, empirical evidence shows a tendency to local optima, if the controller population is started with specialized rules having one weight with  $w = 1$  which would be the equivalent to the full-set-of-parameters approach in the current implementation (data not shown). In turn, the initial uniform distribution of sub-rule weights implies an initial exploration phase as there are many sub-rules that influence the behavior. Later an exploitation phase is characterized by diminishing the weights of counterproductive sub-rules. A more intensive investigation of these interconnections is, however, pending.

## 2.4 Reduction of computational costs

The computation of the hormone dynamics (eq. 3) is costly. In particular, the computational complexity of the summation of the influences of all rules

$$\sum_i \mathcal{S}_i(t) + \mathcal{L}_i(t) + \mathcal{N}_i(t) \quad (13)$$

scales linear with the number of rules  $M$ , that is in Landau notation  $\mathcal{O}(M)$ , and needs to be done in each update step of the controller. This can be avoided by approximating the three functions defined by  $\sum_i \mathcal{S}_i(t)$ ,  $\sum_i \mathcal{L}_i(t)$ , and  $\sum_i \mathcal{N}_i(t)$

for each hormone (and each compartment in case of  $\mathcal{S}$ ) with a stepwise linear function that is pre-calculated and stored in a lookup table. That way only the hormone production, decay, and diffusion needs to be calculated at each time step while the influence through hormones and sensors is determined by simply accessing this lookup table (i.e.,  $\mathcal{O}(1)$ ).

Note the shift of paradigm because now there is a considerable difference between the genotype of the controller, which forms the basis on that it is evolved, and its phenotype, which is the actual representation in the agent at runtime.

A drawback is the increased memory cost that scales exponentially with the number of hormones  $N$  due to the combinatorial possibilities of the hormone-to-hormone sub-rules (i.e.,  $\mathcal{O}(N!)$ ). As we are typically using small numbers of hormones ( $N \leq 3$ ) memory is not an issue and the computational complexity reduces about to that of ANN or even below for the case of using only  $N = 1$  hormone.

The speed-up of using lookup tables instead of computing eq. 3 in each time step is about 50, the speed-up of using AHHS with lookup tables compared to NEAT [63] is about 1.3 (data not shown) in the benchmark reported in the following.

### 3 Domain: Coupled inverted pendulums

The primary domain for AHHS controllers is modular robotics. However, in order to permit high numbers of evolutionary runs we had to restrict our studies to a domain of much lower computational cost. As there is no standard, abstract benchmark for modular robotics yet, that would incorporate typical requirements of modular robotics, we have defined the domain of coupled inverted pendulums [25]<sup>1</sup>.

#### 3.1 Description of the coupled inverted pendulums benchmark

Research of synthesizing controllers for a single inverted pendulum (broom balancing) dates back at least to Widrow and Smith [73]. Applying evolutionary algorithms to this problem dates back at least to Koza and Keane [35]. During the last 20 years the problem was successfully solved which is also true for even more complex scenarios, such as the double pendulum or the triple pendulum (this is however out of the focus of this paper).

We apply several changes to the standard inverted pendulum scenario to increase its complexity and to increase its similarity to challenges of modular robotics scenarios. The pendulums are started in lower positions, that is, we include the nonlinear upswinging phase. We also restrict the cart track length resulting in a scenario similar, for example, to that reported by Chatterjee et al. [10]. In combination with a limited acceleration of the cart motor the upswinging cannot be managed by just moving back and forth once.

---

<sup>1</sup><http://heikohamann.de/coupledInvertedPendulums/>

In addition we limited the sampling rates of all sensors. The sampling rates are low which is documented by the relation between the pre-defined cycle length  $\tau$  of the controller and the maximal angular velocity of  $0.05\pi[1/\tau] = 9^\circ[1/\tau]$ . The pendulum can move up to  $9^\circ$  between two calls of the controller. The controller has little time to adapt to new configurations.

In order to adapt the sensor setting to those that are more typical in robotic scenarios, the sensors do not deliver actual angles and positions directly. These values are partitioned onto several sensors and they are also relative rather than absolute (e.g., distance to wall instead of the cart’s position), see Table 2 for details. All sensor and actuator values are integer values on the interval  $[0, 127]$ . The controllers have two outputs, left actuator  $A_0$  and right actuator  $A_1$  and the acceleration control of the cart is determined by their difference (see Table 2). The most important difference to the standard inverted pendulum is that we couple several carts (or modules) by chains. Carts can move independently as long as they do not pull a chain or run into each other. Hence, each cart has to avoid other carts and walls (cart track ends) and has to balance its pendulum at the same time. Note the difference of this domain from others that mount several pendulums on the same cart, for example, see Xin and Kaneda [77]. This would correspond to a chain length of 0. However, here we are able to define degrees of coupling continuously [25].

In the following experiments we increase the module number (i.e., cart number) without changing the track length. Hence, with increasing module number also the module density is increased which increases the difficulty even more. The modules are controlled locally without global information, ID or positional information and all modules are controlled by identical controllers.

We use an aggregate fitness function [42] which is basically the percentage of time steps that all pendulums spent in the upper equilibrium position ( $\phi = 0$ ). Deviations from  $\phi = 0$  are linearly scaled, that is,  $\phi = 0.5\pi$ , for example, is evaluated as ‘50% in upper position’. A fitness of 1 means all pendulums spent all time in the upper position, 0.5 can be interpreted as ‘the pendulums spent half the time in upper position’, and a fitness of 0 means all pendulums spent all time in the lower equilibrium. If any constraint is violated (e.g., cart runs into other cart, cart runs into chain, cart runs into wall, pendulum velocity too high, etc.) an evaluation run is aborted and the fitness is reduced proportionally to the elapsed time.

The implementation of the cart-pole dynamics is calculated by the Runge-Kutta method of third-order [49] with a discrete time step of size  $\Delta t = 0.01$ .

### 3.2 Discussion of properties of the coupled pendulums benchmark

A detailed discussion of this benchmark is reported in [25]. Search algorithms operating on fitness landscapes related to this benchmark seem to be prone to local optima. Early in an evolutionary run fast motion of the carts earns good fitness improvements. Subsequently further improvements in the fitness can be reached by spinning the pendulums fast. In a third phase of fitness increase the

Table 2: Sensor and actuator setting.

sensor/actuator ID	compartment (AHHS)	sensor/actuator name	mapping of system states to sensor/actuator values	
$S_0$	left	pendulum angle sensor 1	$\phi \in [0, 0.5\pi] \rightarrow [127, 0]$ , 0 else	
$S_1$	right	pendulum angle sensor 2	$\phi \in [\pi, 1.5\pi] \rightarrow [0, 127]$ , 0 else	
$S_2$	left	pendulum angle sensor 3	$\phi \in [0.5\pi, \pi] \rightarrow [127, 0]$ , 0 else	
$S_3$	right	pendulum angle sensor 4	$\phi \in [1.5\pi, 2\pi] \rightarrow [0, 127]$ , 0 else	
$S_4$	left	proximity sensor 1	dist. to obstacle left (max. 1) $\rightarrow [0, 127]$	
$S_5$	right	proximity sensor 2	dist. to obstacle right (max. 1) $\rightarrow [0, 127]$	
$S_6$	left	cart velocity sensor 1	$v \in [-2, 0] \rightarrow [127, 0]$ , 0 else	
$S_7$	right	cart velocity sensor 2	$v \in [0, 2] \rightarrow [0, 127]$ , 0 else	
$S_8$	left	pend. angular vel. sensor 1	$\omega \in [-5\pi, 0] \rightarrow [0, 127]$ , 0 else	
$S_9$	right	pend. angular vel. sensor 2	$\omega \in [0, 5\pi] \rightarrow [0, 127]$ , 0 else	
$A_0$	left	actuator left	$A_0 \in [0, 127]$	$A_0/127 - A_1/127 \rightarrow [-1, 1]$
$A_1$	right	actuator right	$A_1 \in [0, 127]$	




Figure 5: Coupled inverted pendulum benchmark with two carts, pendulums are free to move full  $360^\circ$  mounted on the carts that move in one dimension (left/right) bounded by walls (track ends) and other carts. Marked angle is pendulum angle  $\phi$ .

controller might manage to slow down the pendulums' speed when approaching  $\phi = 0$  but the pendulums still spin. Finally, in the absence of noise an evolved controller could generate deterministic cart trajectories that end up with all pendulums at  $\phi = 0$ . These solutions can be seen as local optima before a fully reactive controller is evolved that actually controls the pendulums also in noisy conditions.

Notice also that conflicting interpretations of sensory inputs exist. A close-by, neighboring cart might be a safe condition in case it moves in the same direction. In contrast a close-by wall might be a dangerous condition.

## 4 Results

In the following, we report several results applying the AHHS approach to the coupled inverted pendulums benchmark. The evolution of AHHS controllers is compared to NEAT [62], we investigate how the performance of AHHS can be improved, several internal processes of AHHS controllers are investigated, and the scaling behavior of AHHS controllers to a higher number of carts is analyzed.

The initial setting as shown in Table 3 (cart and pendulum positions etc.) are maintained except for the explicit investigation of varied initial settings. In case of AHHS we use a mutation rate of 0.4 per hormone and per rule and values are changed at most for  $\pm 60\%$  of the maximal value (randomly uniform). We use a mere linear proportional selection with elitism of 3 (the 3 best individuals are guaranteed to stay in the population).

All of this research is done with a focus on modular robotics applications [24, 67, 51]. The main application is the offline evolution of controllers for modular robots based on simulations. In these simulations the application of a physics engine (simulation of friction, inertia etc.) is important because the evolved behaviors often rely, for example, on friction [75]. A drawback are the very high computational costs that considerably reduce the number of feasible evaluations. Therefore, we have to limit both the population size and the number of

Table 3: Settings of the evolutionary experiments.

min. dist. between carts $d_{\min}$	0.05
max. dist. between carts $d_{\max}$	0.35
track length $w$	2 (on the interval $[-1, 1]$ )
initial pos. $x_i(0)$ of cart $i$	$\mathbf{x}(0) = (-0.4, -0.2, 0, 0.2, 0.4)$
initial pos. $\phi_i(0)$ of pendulum $i$	$\phi(0) = (0.8\pi, 0.9\pi, \pi, 1.1\pi, 1.2\pi)$
initial cart velocities $v_i(0)$	$\mathbf{v}(0) = (0, 0, 0, 0, 0)$
initial pendulum angular vel. $\omega_i(0)$	$\omega(0) = (0, 0, 0, 0, 0)$
population size	100
number of generations	200
mutation rate	0.4
recombination rate	0.01
selection	linear proportional
elitism	3 best individuals

generations. The experiments in this paper are based on the coupled inverted pendulums benchmark, which has low computational costs, and we would have the resources to increase the evaluation number. However, we want to investigate the performance of the AHHS approach with few evaluations in order to transfer the results to the application of offline evolution of controllers for modular robotics. Hence, the population size is limited to 100 genomes and the number of generations is 200 (see Table 3) which results in  $2 \times 10^4$  evaluations.

#### 4.1 Comparison of AHHS to NEAT

We compare the AHHS approach to NEAT [63]. Our implementation of NEAT is based on ‘rtNEAT C++ v1.0.1’ by Kenneth Stanley<sup>2</sup>. The optimization of the parameters used by NEAT is complex. Several settings were tested. Best results were achieved using the parameter settings reported by Whiteson and Stone [72] for the mountain car task. The initial network topology was chosen analogous to the sensor and actuator setting of AHHS except for two extra sensors and two extra actuators that implement two communication channels (one to a possible left neighbor and one to a possible right neighbor). The communication channel is implemented by coupling the output nodes with input nodes of neighboring modules and vice versa. This is necessary to allow similar synchronization possibilities as in case of the AHHS approach through diffusion of hormones. However, experiments did not show significant differences in the performance of NEAT with or without such communication channels (data not shown).

The AHHS controllers are set to  $N = 1$  hormone and  $M = 30$  rules, which proved to be a good setting for this domain. For  $N > 1$  we observed less performance (data not shown) and an analysis of changing rule numbers  $M$  is reported in Sec. 5.1. The modules have a simple, virtual compartmentalization. They have a left and a right compartment and the sole hormone  $H_0$  occurs in

<sup>2</sup><http://www.cs.utexas.edu/users/nn/keyword?rtneat>




Figure 6: Box-and-whisker plots of the fitnesses of the best evolved controllers for the performance comparison between AHHS and NEAT with several module numbers and  $n = 30$  runs each. Asterisks show significances of  $p < 0.05$  using the Wilcoxon rank-sum test.

two concentrations  $H_0^0$  (left) and  $H_0^1$  (right) accordingly. The hormone diffuses between the two compartments of the module as well as between neighboring compartments of neighboring modules with a defined diffusion coefficient  $D_0$  (cf. eq. 3, see [24] for a more detailed discussion of module-to-module diffusion, for example Fig. 11b therein). Sensors and actuators are also associated with one of the two compartments (see Table 2). The population is initialized with random controllers (uniformly distributed sub-rule weights, fixed and dependent doses etc.).

The comparison was done for 1 through 5 modules with  $n = 30$  evolutionary runs each. The results are shown in Fig. 6. For the case of one module NEAT outperforms AHHS significantly. However, for all multi-module settings AHHS significantly outperforms NEAT. From these results we follow that AHHS is at least competitive to other state-of-the-art approaches in case of decentrally controlled multi-module systems.

## 4.2 Applying noise during evolution for better performance

In the following experiment we have applied noise to the hormone during the evolution. A uniformly random distributed noise on the interval  $[-0.05, 0.05]$  was added to the hormone concentrations during generations  $[0, 30]$ ,  $[60, 100]$ , and  $[140, 160]$ . During these generations the fitness was based on the minimum




Figure 7: Box-and-whisker plots of the fitnesses of the best evolved controllers for the method that applies noise to the hormone compared to the performance of the standard AHHS method as shown in Fig. 6.

fitness out of 3 evaluations. The idea of introducing phases of noisy hormone concentrations is to have an additional exploration in the search space and to find more robust solutions. The phases without noise are important to find the best controllers for the deterministic domain which is our benchmark here. The results are shown in Fig. 7, the performance achieved by applying noise is compared to the AHHS performance shown in Fig. 6. Significantly better performance was achieved for 5 modules, while significantly worse performance was achieved for 1 module. For 2, 3, and 4 modules only a trend seems to indicate better performance by applying noise. Hence, it is questionable whether the higher computational costs due to the higher number of evaluations are justified.

### 4.3 Analysis of an evolved controller

In the following we investigate the best evolved controller that was achieved using the noise method described above for the 2 module scenario. The behavior of the evolved controller in a regular evaluation run is shown in Fig. 8. For this initial state it achieves a fitness of about 0.85 which is equivalent to keeping both pendulums upwards balanced for 3400 time steps (the actual value is smaller because during the up-swinging phase pendulum angles of  $\phi \neq 0$  occur). Both carts show only small movements and move in synchrony which is obviously a good behavior to avoid collisions. The carts keep moving towards  $x = -1$  which



indicates that the pendulums probably won't be balanced for much longer than the evaluation period. However, keeping the carts around one position was not part of the fitness function. Note the complex dynamics of the hormone which is an oscillation over 3 time steps of the form:  $0, 1, x, 0, 1, x, \dots$ . Hormone concentrations of 0 and 1 correspond to actuator control values of 0. Hence, the actuator actually is active only every third time step.

One great benefit of the AHHS approach is that a complete qualitative representation of the evolved controllers is possible as shown in Fig. 9. This representation is based on the lookup tables, that map sensory input to hormone concentrations, hormone concentrations to hormone concentrations and hormone concentrations to actuator control values, and can be generated automatically. In this graphical description of the controller it is, for example, easy to see that sensors  $S_0$  to  $S_6$  and actuator  $A_1$  are basically turned off. Hence, no collision avoidance behavior was evolved because input from the proximity sensors is ignored. The sensors of the pendulum angles are also not used.

The internal dynamics (i.e., without dynamic sensory input) is determined by the feedback loop through hormone-to-hormone rules. A bifurcation diagram of the hormone-to-hormone mapping function as shown in Fig. 9 is shown in Fig. 10. In order to simplify the dynamics of the complete system, the overall influence by sensors  $\sum_i \mathcal{S}_i = S_c$  is assumed to be constant and is added to the hormone concentration in every time step. That yields the following map for the hormone dynamics

$$H(t+1) = H(t) + \alpha H(t) - \mu H(t) + S_c + \sum_i \mathcal{L}_i + \mathcal{N}_i \quad (14)$$

In Fig. 10 the constant  $S_c = \sum_i \mathcal{S}_i$  is used as bifurcation parameter. For each value  $\sum_i \mathcal{S}_i$  there are infinitely many combinations of associated sensor values that are not individually considered. The diagram shows 3 qualitatively different regions: a period-2 region ( $\sum_i \mathcal{S}_i \in [-1, -0.229]$ ) a period-3 region ( $\sum_i \mathcal{S}_i \in [-0.229, 0.329]$ ), and a fixed point region ( $\sum_i \mathcal{S}_i \in [0.329, 1]$ ). The controller mainly operates on the interval of approximately  $\sum_i \mathcal{S}_i \in [0, 0.3]$  (data not shown), which is the period-3 region (i.e., attractors based on oscillations between 3 different hormone concentrations with a period length of 3 time steps).

## 5 AHHS – Mode of operation

In the following we investigate the mode of operation of AHHS controllers. We test the evolved selection of sensors and the dependence on sensory input. The hormone dynamics is interpreted in a nonlinear dynamics context in extension to the results shown in Fig. 10. Finally, we investigate the emergence of oscillations and different time-scales in the hormone dynamics.

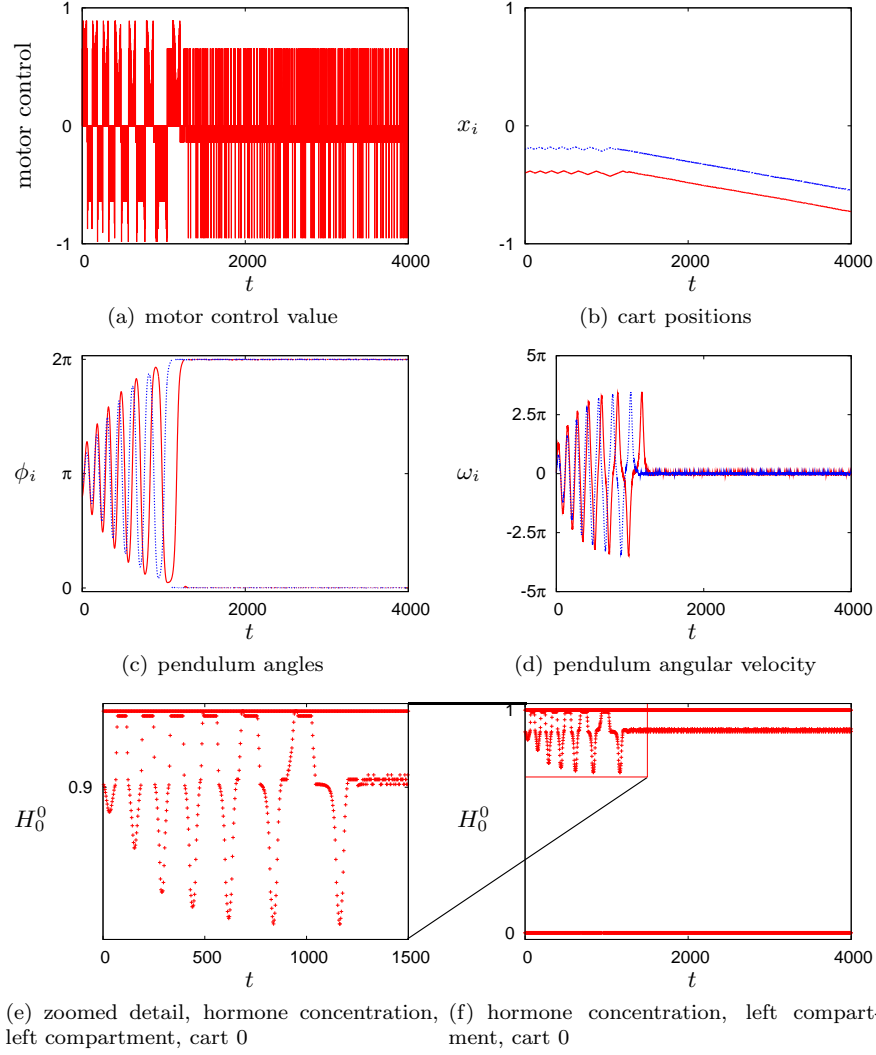


Figure 8: System states for a run of an evolved controller for 2 modules. Note the complex dynamics of the hormone which is an oscillation of period 3 time steps, cf. Fig. 9.

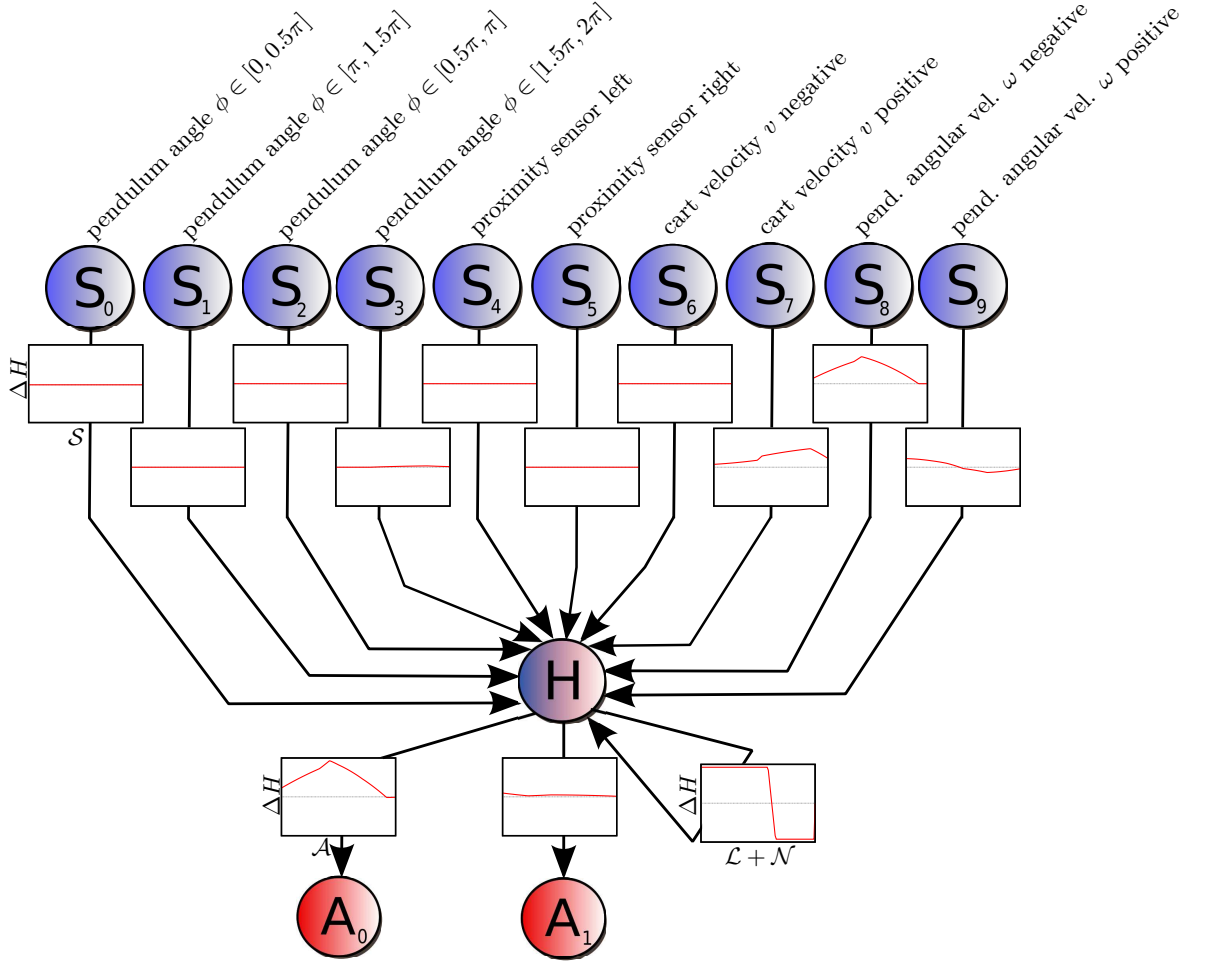


Figure 9: Complete qualitative representation of the best evolved AHHS controller. The upper row of nodes represents the 10 sensors. The 10 small diagrams show the sensor-to-hormone mappings. The node in the center represents the hormone concentration and the neighboring arrow loop represents the hormone-to-hormone mapping. The 2 lower nodes represent the actuator outputs with the corresponding hormone-to-actuator mappings.




Figure 10: Bifurcation diagram for the internal hormone concentration dynamics of the evolved controller shown in Fig. 9. The overall influence by sensory input  $\mathcal{S}$  is used as bifurcation parameter (see eq. 14). The controller mainly operates in the period-3 region of about  $\sum_i \mathcal{S}_i \in [0, 0.3]$  (cf. dynamics of the hormone concentration in Fig. 8(f)).

### 5.1 Influence of rule numbers

We begin by investigating the influence of allowed rules  $M$  for the setting with 2 modules. Results are shown in Fig. 11. On the one hand, a certain number of rules (from Fig. 11 one can tell that  $M = 30$  would most likely be a good choice here) is necessary to reach good performance. On the other hand, high rule numbers seem not to decrease performance relevantly. Note that the dependence on the rule numbers is not simply a question of whether a function (as shown in Fig. 9) is representable by few rules but rather how much of the space of all possible behaviors is covered by the initial population. Hence, a well chosen rule number would maximize the heterogeneity of behavior in the initial population without increasing the search space unnecessarily. This relation, however, needs more investigation and will be future work together with the option of evolving the number of rules also (i.e., dynamic genome lengths).

### 5.2 Most influential sensors

The approach with just one hormone proposed in this work allows a straight forward time series analysis of the internal states (described by a single hormone concentration) compared with the sensory input. We use Pearson’s product-moment correlation (PMCC) for the temporal differences (first derivatives) for pairs of the hormone concentration and sensor values. By this investigation we want to detect the sensors that mainly influence the controllers. This way we regard the evolved controllers as black boxes and analyze their behavior statistically.




Figure 11: Box-and-whisker plots of the fitnesses of the best evolved controllers for several rule numbers  $M \in \{2, 4, 6, 10, 20, 30, 40, 50, 60, 70, 80, 90, 100, 200\}$  (significances not displayed).

The correlation coefficient  $r$  gives a good estimate on how relevant the particular sensor is. Values of  $r = 1$  or  $r = -1$  indicate a positively linear or negatively linear relationship. Values of  $|r| < 1$  indicate an increasing noise until maximum noise is reached for  $r = 0$ . Note that nonlinear relationships cannot be discovered with PMCC.

Results for the best controllers, that were evolved applying the noise method (see Fig. 7), are shown in Fig. 12. Interestingly no controller listens mainly for the pendulum angles (sensors 1 through 4) and those controllers listening for the proximity sensors correlate almost always positively. The best evolved controllers listen only to sensors 6 through 9 being cart velocity and pendulum velocity. These sensors are commonly regarded as the most relevant in inverted pendulum scenarios [46].

### 5.3 Nonlinear dynamics interpretation

The bifurcation diagram shown above (Fig. 10) already indicated that AHHS controllers can be investigated by applying methods of nonlinear dynamics. The hormone-to-hormone sub-rules define the internal dynamics by a feedback loop. This internal dynamics is then disturbed by sensory input. The dynamics of the hormone concentrations can be interpreted as trajectories in a state space. A demonstration (based on a controller with several hormones in contrast to those reported above) in form of a vector field of such a state space focusing on one hormone-to-hormone relationship is shown in Fig. 13(a). The vector field has a fixed point at about (0.64, 0.93). The occurring states of a regular evaluation run shown in Fig. 13(a) indicate that the system, however, stays almost never at the fixed point. In fact, sensory input and possibly also the





Figure 12: Correlation coefficients for correlations between changes of hormone value (first derivative) and the changes of the ten sensors in the left and right compartment; dashed lines show thresholds of high significance ( $p < 0.01$ ) for this data set (outside the two lines all data is significant, inside only two data points).

influence of other hormones disturb the system and drive it temporally away from the fixed point.


Fig. 13(b) shows the hormone concentration  $H_0$ , the most influential sensor  $S_9$  (negative angular velocities), and the corresponding pendulum angular velocity  $\omega$  for the initial 2500 time steps. Without any disturbance by sensory input, hormone 0 would relax to the steady state of about  $H_0 = 0.93$ . The sensory input  $S_9$  adds to the hormone and increases the hormone concentration to  $H_0 = 1$ . Without input from  $S_9$  the hormone oscillates irregularly between about 0.75 and 1 due to other inputs. The actuator is activated in these phases resulting in short periods of accelerations of the cart that keep the pendulum swinging.

#### 5.4 Control based on oscillations

An interesting property of all evolved controllers is the oscillation of all hormone concentrations. Certainly the domain of oscillating pendulums introduces oscillations to the system but the controllers evolved also much faster oscillations. Even fast swinging pendulums typically have periods of more than 200 time steps whereas hormones often show periods of 2 to 7 time steps. Obviously controllers whose functionality is based on such oscillations, are favored by the evolutionary algorithm in this domain. In fact, these oscillations can be interpreted as means of storing system states. This is clearly indicated by Fig. 14 that shows a hormone with two different kinds of oscillation behavior. Hence, one can argue that in addition to those oscillations of the pendulum a second time-scale emerges. The emergence of time-scales and their significance as forms of memory and for robustness were reported regularly [20, 36, 52]. Similar dynamics were found in single neurons [15]. In this context, note also the



(a) vector field and trajectory of  $(H_0, H_1)$



(b) hormone concentration, sensor, and pendulum dynamics

Figure 13: Hormone dynamics: vector field as defined by the linear and non-linear hormone-to-hormone rules for two selected hormones, dynamics of these hormones, corresponding sensory input and angular velocity of the pendulum of one module.




Figure 14: Controller states represented by different oscillations; hormone dynamics and angle of the pendulum of one module in a two-pendulum scenario.

discussion of Moiola et al. [38] in Section 8. Considering these fast oscillations, the similarity to neurons, and our inspiration by hormones one could argue that there is a contradiction. However, we argue that the dynamics of hormones in unicellular organisms is also fast compared to their environments (e.g., nutrition gradients).

## 6 Generalization and Scaling

All controllers reported in this paper were evolved for only one initialization of module numbers and cart and pendulum positions. In the following we investigate how these controllers generalize to other initializations and to other module numbers. Both are desirable because it would allow to evolve the controllers in scenarios of lower complexity (e.g., only 3 modules and only one initialization of cart and pendulum positions) than the final application scenario which is resource-conserving.

### 6.1 Generalization to initial conditions far from original domain

In the following we investigate how evolved controllers generalize to initial angles of pendulums for which the controllers were not evolved. This was done for the best controllers of the 2-module scenario using the AHHS-noise approach that is shown in Fig. 7. Without re-evolving the controllers they were tested for all pairs of initial pendulum angles  $(\phi_0(0), \phi_1(0))$  on the full interval  $\phi_i(0) \in [0, 2\pi]$ , for  $i \in \{0, 1\}$  with a resolution of  $\Delta\phi = 0.05$ .



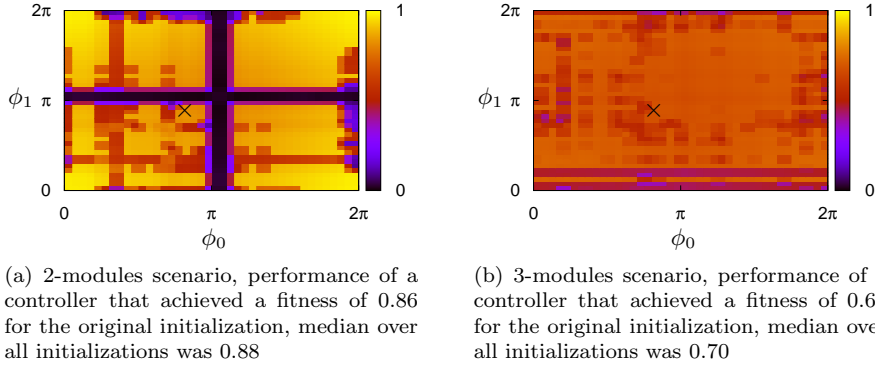


Figure 15: Fitness for 1681 initial pendulum angles of the first 2 pendulums  $\phi_0$  and  $\phi_1$ . The crosses at  $(0.8\pi, 0.9\pi)$  mark the initialization for which the controller was evolved (in case of 3-module scenario we have additionally  $\phi_2(0) = \pi$ ).

The median fitness over all 30 controllers and all tested initial angles was 0.61 (1st quartile: 0.23, 3rd quartile: 0.70). This is a good result compared to the NEAT experiments, that were obtained by explicitly evolving for one initialization and reached only a median of 0.59 (see Fig. 6).

In Fig. 15(a) the results for a controller are shown that achieved a fitness of 0.86 for the original initialization. The median over all initializations for this controller was 0.88 (1st quartile: 0.28, 3rd quartile: 0.92) which is above the fitness for the original initialization. The crosswise black shape in Fig. 15(a) is representable for 28 of the 30 evaluated controllers and indicates low fitness for initializations of  $\phi_i(0) \approx \pi$  which is about the lower equilibrium of the pendulum. These initializations are difficult for the controllers because they are not part of the original initialization of the 2-module scenario (in contrast to the 3, 4 and 5 module scenarios, see Table 3). For  $\phi = \pi$  the pendulum is at rest and the important sensory stimulus from the angular velocity sensors is missing. In similar experiments for the 3-module scenario (i.e., controllers evolved in the 3-module scenario were evaluated on the full interval of  $\phi_i$ ) the crosswise shapes vanish (see Fig. 15(b)) because the initialization  $\phi_i(0) = \pi$  is part of the initialization for which the controllers are evolved.

Similar results were reported before for ANN and a single inverse pendulum by Pasemann [46]. Adaptive controllers, that worked well also on initial conditions far from the original domain, emerged for the single pendulum.

## 6.2 Scaling of module numbers

In the following we investigate how evolved controllers cope with module numbers that are higher and lower than in the scenario they have been evolved for. The respectively best controller of the 30 runs for the 3-module scenario using the AHHS-noise approach were evaluated in 1-, 2-, 4-, and 5-module scenarios.




Figure 16: Scaling of module numbers; the controllers were evolved for the 3-module setting (data shown in Fig. 7). The same controllers were evaluated in the 1-, 2-, 4-, and 5-module setting without re-evolving them. There is no significant difference in the performance.

The result is shown in Fig. 16. There is no significant difference between the 5 settings. Hence, the controllers scale up and down equally well. Effective scaling opens up the possibility of evolving controllers for simpler scenarios that consume less resources. Later the evolved controller is applied to more complex scenarios. The feasibility of this approach is also supported by the good scaling property in individual cases of controllers that were evolved for the 2-module scenario and that showed high fitness in the 5-module scenario close to the all-time record.

## 7 Comparison to similar implementations

In order to test the relevance of defining the sensor-to-hormone, hormone-to-hormone, and the hormone-to-actuator mappings using the AHHS concept of rules, we have implemented two reference implementations. We call the first one ‘direct table approach’ (DT). It uses the lookup tables of the AHHS implementation but the table entries are filled directly without the concept of summations of functions defined by sub-rules. Initially the tables are filled with random numbers that are directly mutated by the genetic algorithm. Neighboring values in the table are not correlated in contrast to the AHHS approach.

The second reference implementation is called ‘Fourier approach’. It implements functions using Fourier series to fill the lookup tables. This way neighboring values in the table are correlated because Fourier series are differentiable which should also be reflected in discrete samplings of Fourier series. The Fourier coefficients are initialized randomly and mutated by the genetic algorithm.




Figure 17: Box-and-whisker plots of the fitnesses of the best evolved controllers for the performance comparison between the direct encoding of lookup tables (DT) and the best AHHS approach respectively (AHHS-noise approach for several modules and the standard approach for one module).

## 7.1 Direct encoding of lookup tables

This reference implementation uses the lookup tables of the AHHS implementation but the entries are filled directly. Note that this preserves the general network topology as shown in Fig. 9. However, instead of filling the tables by the summation of several sub-rules the tables are filled directly. Initially they are filled with random numbers. The mutations act on these tables directly, too. The following results were obtained with tables of 128 bins (i.e., the intervals of sensor values and the intervals of hormone concentrations are discretized with a resolution of 128 discrete steps) and with an average of 320 mutations per controller and generation. The results are shown in Fig. 17. The performance of the direct table approach seems to be independent of the module number. This is explained by the local optimum, that is typically found, which is a fast pendulum swinging behavior without any interaction between the carts. Besides the quantitatively small difference of the DT and AHHS medians of the 4- and 5-module scenario there is a qualitative difference because the DT-behaviors show only fast pendulum swinging while the AHHS-behaviors use more promising strategies (e.g., a single pendulum is balanced).

In addition, this simple approach shows significantly better performance for 3 and 5 modules compared to NEAT. This might be seen as an indicator that shaping sensor-to-hormone and hormone-to-actuator mappings is a generally beneficial approach.

## 7.2 Lookup tables based on Fourier series

In this second reference implementation we use Fourier series of order  $k = 24$ . Such a high order might be counter-intuitive but proved empirically to be a better approach than using lower orders. Obviously short wavelengths in the functions are advantageous. The lookup tables are filled by functions of the form

$$f(x) = \frac{a_0}{2} + \sum_{k=1}^{24} (a_k \cos(kx) + b_k \sin(kx)), \quad (15)$$

with  $x \in [0, 2\pi]$  representing the scaled sensor/hormone value associated with the considered table entry. The values  $a_0, a_k, b_k$  ( $k \in [1, 24]$ ) are initialized randomly and mutated by the genetic algorithm.

The results shown in Fig. 18 indicate that the rule-based AHHS approach reported in this paper is not the general single best approach for all scenarios. The AHHS approach reaches significantly better performance for the 1- and 2-module scenarios. However, the Fourier series approach has higher computational cost. Our implementation ran slower by a factor of about 2 compared to the AHHS-noise approach reported in Sec. 4.2. Reducing the order of the Fourier series would lower the computational cost but also the performance. Hence we follow that the AHHS approach is still the better choice for this domain.

## 8 Related work

The proposed AHHS controller approach has methodological similarities to many other approaches, however, the unrestricted evolution of hormone-to-hormone reactions as well as sensor-to-hormone and hormone-to-actuator mappings seem to be unique features. In the following we compare our approach to neural network approaches, gene regulatory networks, and reaction-diffusion approaches.

### 8.1 Artificial neural networks, NEAT, HyperNEAT, and GasNet

It seems to be justified to call ANN (in particular, continuous-time recurrent neural networks) the standard approach in the automatic synthesis of controllers, especially in the field of evolutionary robotics [43, 44, 14, 31, 27, 69]. However, the standard ANN approach fails for complex tasks which is, as mentioned in the introduction, documented by the absence of complex tasks in the literature [42]. The difference of the AHHS approach to ANN is mainly the concept of hormone diffusion and, as mentioned above, the X-to-hormone mappings that can be interpreted as network functions with weights depending on the outputs of connected neurons.

In contrast to the standard ANN approach, NEAT [63] and HyperNEAT [64] have verifiably higher evolvability. NEAT differs in the sophisticated methods




Figure 18: Box-and-whisker plots of the fitnesses of the best evolved controllers for the performance comparison between the approach based on Fourier series and the best AHHS approach respectively (AHHS-noise approach for multiple modules and the standard approach for one module).

of evolutionary computation while the underlying ANN is that of the standard approach.

AHHS can be compared to ANN in the following way: Assume (and for one example it is empirically shown in this paper) that it is a good way to solve the problem of controller synthesis by the form of function fitting for sensor-to-hormone, hormone-to-hormone, and hormone-to-actuator mappings as reported in this work. Obviously ANN are general enough to solve it principally in the same way. However, it is improbable that separated function-fitting clusters, that implement these mappings, emerge themselves within the network, even if NEAT is used.

HyperNEAT differs to NEAT by its generation of connectivity patterns that could be compared to the spatiality introduced by the compartmentalization using AHHS. Otherwise there seem not to be many similarities between these two approaches.

Much more similarities are found in comparing AHHS to GasNet [31] because in addition to the standard ANN they introduce diffusing hormones and, hence, represent hybrid approaches. The related artificial homeostatic systems [70, 71, 41] emphasize the homeostatic aspect similar to the internal dynamics of AHHS.

## 8.2 Gene regulatory networks and reaction-diffusion systems

Most similarities are found by comparing AHHS to other approaches that do not rely on ANN. Gene regulatory networks (GRN) [3, 18] that are inspired by biological gene networks implementing an implicit encoding of networks. Still there are many differences between GRN and AHHS although superficially there seem to be a lot of similarities due to the similar biological inspiration. A difference in the genotypes is that AHHS do not have the distinction between coding/non-coding regions. The AHHS genome is rather predefined and static. Concerning the resulting networks and their functioning, each edge has its own activation threshold and redundant edges with different activations are allowed in the AHHS approach.

Probably most similarities are found in the comparison to the diffusion-reaction controller approach [13]. Similarities are the definition of controllers based on reactions and the essential influence of diffusion processes. Differences are the predefined reactions (Gray-Scott reaction-diffusion system) in Dale’s approach. In addition, spatiality is playing a much more important role as they use 128 compartments while our agents have only two compartments in this work. Dale’s approach relies on this high resolution and, hence, shows spatial qualities that are not observed in the AHHS approach.

Furthermore, similarities are found in comparison to the controller based on the Kuramoto model by Moiola et al. [38]. In that work, the focus is on synchronization based on predefined oscillations. In the AHHS approach, oscillations and synchronization emerge themselves.

The evolution of sensor-to-hormone mappings can be interpreted as a direct function fitting approach related to Jang [32]. However, the AHHS approach is

much more sophisticated and defines these mappings not directly but in a form that shows high evolvability.

Other approaches that are inspired by hormones are very different from ours. For example, Shen et al. [59] report a message-based system and Shen et al. [60] report a pheromone-like implementation (comparable to Payton et al. [47]). Furthermore, see Meng et al. [37] for a modular robotics approach inspired by a mechanochemical model for cell morphogenesis—an aspect that we do not address here but elsewhere, see Schmickl et al. [55].

## 9 Conclusion

In this paper we have reported an application and intensive analysis of the AHHS controller approach, which is inspired by signaling networks in unicellular organisms. We have reported a comparison to NEAT in the coupled inverted pendulums benchmark that revealed a significantly better performance of the AHHS controller for multi-module settings with a limited number of evaluations. It can be considered a main result of this work that the AHHS approach can be interpreted as an ANN with a rather exceptional network function based on weight functions depending on the outputs of connected neurons. This concept might also be a promising approach for ANN. The lookup-table implementation of the AHHS approach introduces an interesting difference between genotype and phenotype of the controllers. The genotype is designed for high evolvability while the phenotype is optimized for low computational complexity.

An advantage of AHHS compared, for example to ANN, NEAT or HyperNEAT, is that a full but simple representation of an AHHS controller is possible as shown in Fig. 9. Note that the networks evolved by the NEAT approach are much more complex and contain between 25 and 40 nodes.

A nonlinear dynamics analysis of the evolved controllers is possible, especially due to the one-dimensionality as we are using only one hormone in this work. This analysis showed that the controllers are mostly based on periodic attractors that emerge during evolution.

The investigation of two reference implementations showed that evolving X-to-hormone mappings based on summations of functions is advantageous to direct encodings of such mappings. Hence, our approach of evolving the mappings indirectly via groups of tunable sub-rules seems to be particularly suitable for methods of artificial evolution.

Generally the AHHS approach does not require predefined topologies. In this work the network topology was not predefined because the controllers were initialized with fully connected networks. Unnecessary connections had to be turned off explicitly by evolution as it is seen in Fig. 9 (zero lines for sensors 0 through 6). The probability of fully connected networks is reduced either by decreasing the number of rules  $M$  or by increasing the number of hormones  $N$ . The relation between the rule number and the hormone number changes also the representable shapes of the X-to-hormone mappings (cf. Fig. 4(a)). The initialization with almost no connections would also be possible, if the addition

of rules (i.e., growing genome lengths) would be allowed.

The application of the AHHS approach to other domains has been shown previously. The evolution of controllers for single robots and for modular robotics has been shown in [66, 23, 24]. The analysis of single robot controllers was reported in [54]. The implementation of a simple AHHS controller on hardware was reported in [65, 56].

The analysis has shown that the proposed approach is superior to state-of-the-art approaches in the considered domain with local control, evolved controllers generalize to other initializations and scale with the number of modules. This opens up many possibilities. For example, controllers can be evolved within simpler scenarios in order to save resources. Furthermore, also changes in the number of coupled modules could be managed dynamically by the same controller without external adaptations. The good scaling and generalization properties could probably even be improved by re-evolving the controllers in new or more complex domains. In addition, this method could be combined with environmental incremental evolution as reported by Nakamura et al. [40]. Our future work will mainly be the application of the proposed approach to the multi-modular robotics domain.

## Acknowledgment

We thank the anonymous referees for helpful comments. This work is supported by: EU-IST-FET project ‘SYMBRION’, no. 216342 and by EU-ICT project ‘REPLICATOR’, no. 216240.

## References

- [1] B. Alberts. *Molecular biology of the cell*. Garland Pub., 1989.
- [2] H. Armus, A. Montgomery, and J. Jellison. Discrimination learning in *Paramecia* (*p. caudatum*). *The Psychological Record*, 56:489–498, 2006.
- [3] J. C. Bongard. Evolving modular genetic regulatory networks. In *Proceedings of the 2002 Congress on Evolutionary Computation (CEC’02)*, pages 17–21, 2002.
- [4] D. Bray. Intracellular signalling as a parallel distributed process. *Journal of theoretical biology*, 143(2):215–231, 1990. ISSN 0022-5193.
- [5] D. Bray. Protein molecules as computational elements in living cells. *Nature*, 376:307–312, 1995.
- [6] D. Bray. *Wetware: A Computer in Every Living Cell*. Yale University Press, 2009.
- [7] R. Brooks. A robust layered control system for a mobile robot. *IEEE Journal of Robotics and Automation*, 2(1):14–23, 1986.



- [8] R. Brooks. From earwigs to humans. *Robotics and Autonomous Systems*, 20(2-4):291–304, 1997.
- [9] S. Camazine, J.-L. Deneubourg, N. R. Franks, J. Sneyd, G. Theraulaz, and E. Bonabeau. *Self-Organization in Biological Systems*. Princeton University Press, 2001. ISBN 978-0-691-11624-2.
- [10] D. Chatterjee, A. Patra, and H. K. Joglekar. Swing-up and stabilization of a cart-pendulum system under restricted cart track length. *Systems & Control Letters*, 47(4):355–364, 2002.
- [11] T. Chouard. Revenge of the hopeful monster. *Nature*, 463:864–867, February 2010.
- [12] D. Cliff, I. Harvey, and P. Husbands. Explorations in evolutionary robotics. *Adaptive Behavior*, 2, 1993.
- [13] K. Dale and P. Husbands. The evolution of reaction-diffusion controllers for minimally cognitive agents. *Artificial Life*, 16(1):1–19, 2010.
- [14] M. Dorigo, V. Trianni, E. Şahin, R. Groß, T. H. Labella, G. Baldassarre, S. Nolfi, J.-L. Deneubourg, F. Mondada, D. Floreano, and L. M. Gambardella. Evolving self-organizing behaviors for a swarm-bot. *Autonomous Robots*, 17(2-3):223–245, 2004.
- [15] A. V. Egorov, B. N. Hamam, E. F. M. E. Hasselmo, and A. A. Alonso. Graded persistent activity in entorhinal cortex neurons. *Nature*, 420:173–178, November 2002.
- [16] J. A. Endler. Some general comments on the evolution and design of animal communication systems. *Philosophical Transactions of the Royal Society B*, 340:215–225, 1993.
- [17] D. Floreano, S. Mitri, S. Magnenat, and L. Keller. Evolutionary conditions for the emergence of communication in robots. *Current Biology*, 17:514–519, 2007.
- [18] D. Floreano, P. Dürri, and C. Mattiussi. Neuroevolution: From architectures to learning. *Evolutionary Intelligence*, 1:47–62, 2008.
- [19] D. Floreano, P. Husbands, and S. Nolfi. Evolutionary robotics. In B. Siciliano and K. Oussama, editors, *Handbook of Robotics*, pages 1423–1452. Springer-Verlag, Berlin, 2008.
- [20] K. Fujimoto and K. Kaneko. How fast elements can affect slow dynamics. *Physica D: Nonlinear Phenomena*, 180(1-2):1–16, 2003.
- [21] W. Grey Walter. An imitation of life. *Scientific American*, 182(5):42–45, 1950.

- [22] W. Grey Walter. A machine that learns. *Scientific American*, 185(2):60–63, 1951.
- [23] H. Hamann, J. Stradner, T. Schmickl, and K. Crailsheim. A hormone-based controller for evolutionary multi-modular robotics: From single modules to gait learning. In *Proceedings of the IEEE Congress on Evolutionary Computation (CEC’10)*, pages 244–251, 2010.
- [24] H. Hamann, J. Stradner, T. Schmickl, and K. Crailsheim. Artificial hormone reaction networks: Towards higher evolvability in evolutionary multi-modular robotics. In H. Fellersmann, M. Dörr, M. M. Hanczyc, L. L. Laursen, S. Maurer, D. Merkle, P.-A. Monnard, K. Støy, and S. Rasmussen, editors, *Proc. of the ALife XII Conference*, pages 773–780. MIT Press, 2010.
- [25] H. Hamann, T. Schmickl, and K. Crailsheim. Coupled inverted pendulums: a benchmark for evolving decentral controllers in modular robotics. In N. Krasnogor and P. L. Lanzi, editors, *Proceedings of the 13th Annual Genetic and Evolutionary Computation Conference, GECCO 2011*, pages 195–202. ACM, 2011. ISBN 978-1-4503-0557-0. URL <http://doi.acm.org/10.1145/2001576.2001604>.
- [26] W. Hand and W. Haupt. Flagellar activity of the colony members of *Volvox aureus* ehrbg. during light stimulation. *Journal of Eukaryotic Microbiology*, 18(3):361–364, 1971.
- [27] I. Harvey, P. Husbands, D. Cliff, A. Thompson, and N. Jakobi. Evolutionary robotics: the sussex approach. *Robotics and Autonomous Systems*, 20(2-4):205–224, 1997.
- [28] M. D. Hauser. *The Evolution of Communication*. MIT Press, Cambridge, MA, 1996.
- [29] J. H. Holland. *Adaptation in Natural and Artificial Systems*. Univ. Michigan Press, Ann Arbor, MI, 1975.
- [30] S. Holmes. Phototaxis in *volvox*. *Biological Bulletin*, 4(6):319–326, 1903.
- [31] P. Husbands. Evolving robot behaviours with diffusing gas networks. In *Evolutionary Robotics*, volume 1468/1998 of *Lecture Notes in Computer Science*, pages 71–86. Springer Berlin / Heidelberg, 1998.
- [32] J.-S. R. Jang. Self-learning fuzzy controllers based on temporal backpropagation. *IEEE Transactions on Neural Networks*, 3(5):714–723, 2002.
- [33] S. A. Kauffman and S. Levin. Towards a general theory of adaptive walks on rugged landscapes. *Journal of Theoretical Biology*, 128(1):11–45, 1987.
- [34] J. Koza. *Genetic Programming: On the Programming of Computers by Means of Natural Selection*. MIT Press, 1992. ISBN 0-262-11170-5.

- [35] J. R. Koza and M. A. Keane. Genetic breeding of non-linear optimal control strategies for broom balancing. In A. Bensoussan and J. Lions, editors, *Analysis and Optimization of Systems*, volume 144 of *LNCS*, pages 47–56. Springer-Verlag, 1990.
- [36] A. Kremling, S. Fischer, T. Sauter, K. Bettenbrock, and E. Gilles. Time hierarchies in the escherichia coli carbohydrate uptake and metabolism. *BioSystems*, 73(1):57–71, 2004.
- [37] Y. Meng, Y. Zhang, and Y. Jin. Autonomous self-reconfiguration of modular robots by evolving a hierarchical mechanochemical model. *IEEE Computational Intelligence Magazine*, 6(1):43–54, February 2011.
- [38] R. Moiola, P. A. Vargas, and P. Husbands. Exploring the kuramoto model of coupled oscillators in minimally cognitive evolutionary robotics tasks. In *WCCI 2010 IEEE World Congress on Computational Intelligence - CEC IEEE*, pages 2483–2490, 2010.
- [39] S. Murata, K. Kakomura, and H. Kurokawa. Toward a scalable modular robotic system - navigation, docking, and integration of m-tran. *IEEE Robotics & Automation Magazine*, 14(4):56–63, 2008.
- [40] H. Nakamura, A. Ishiguro, and Y. Uchikawa. Evolutionary construction of behavior arbitration mechanisms based on dynamically-rearranging neural networks. In *Proceedings of the 2000 Congress on Evolutionary Computation*, volume 1, pages 158–165. IEEE, 2000. doi: 10.1109/CEC.2000.870290.
- [41] M. Neal and J. Timmis. Timidity: A useful mechanism for robot control? *Informatica*, 4(27):197–204, 2003.
- [42] A. L. Nelson, G. J. Barlow, and L. Doitsidis. Fitness functions in evolutionary robotics: A survey and analysis. *Robotics and Autonomous Systems*, 57:345–370, 2009.
- [43] S. Nolfi and D. Floreano. Learning and evolution. *Autonomous Robots*, 7: 89–113, 1999.
- [44] S. Nolfi and D. Floreano. *Evolutionary Robotics: The Biology, Intelligence, and Technology of Self-Organizing Machines*. MIT Press, 2004.
- [45] B. Østman, A. Hintze, and C. Adami. Critical properties of complex fitness landscapes. In H. Föllmann, M. Dörr, M. M. Hanczyc, L. L. Laursen, S. Maurer, D. Merkle, P.-A. Monnard, K. Støy, and S. Rasmussen, editors, *Proc. of the ALife XII Conference*, pages 126–132. MIT Press, 2010.
- [46] F. Pasemann. Evolving neurocontrollers for balancing an inverted pendulum. *Network: Computation in Neural Systems*, 9(4):495–511, 1998.
- [47] D. Payton, M. Daily, R. Estowski, M. Howard, and C. Lee. Pheromone robotics. *Autonomous Robots*, 11(3):319–324, Nov. 2001.

- [48] D. Peak, J. D. West, S. M. Messinger, and K. A. Mott. Evidence for complex, collective dynamics and emergent, distributed computation in plants. *Proceedings of the National Academy of Science*, 101(4):918–922, 2004.
- [49] W. H. Press, S. A. Teukolsky, W. T. Vetterling, and B. P. Flannery. *Numerical Recipes in C++*. Cambridge Univ. Press, 2002.
- [50] I. Rechenberg. *Evolutionssstrategie '94*. Frommann Holzboog, 1994. ISBN 978-3-7728-1642-0.
- [51] REPLICATOR. Project website, 2011. <http://www.replicators.eu>.
- [52] I. Rojdestvenski, M. Cottam, Y.-I. Park, and G. Öquist. Robustness and time-scale hierarchy in biological systems. *BioSystems*, 50(1):71–82, 1999.
- [53] M. Rubenstein and W.-M. Shen. Scalable self-assembly and self-repair in a collective of robots. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), St. Louis, Missouri, USA*, St. Louis, Missouri, USA, Oct 2009.
- [54] T. Schmickl and K. Crailsheim. Modelling a hormone-based robot controller. In *MATHMOD 2009 - 6th Vienna International Conference on Mathematical Modelling*, 2009.
- [55] T. Schmickl, H. Hamann, J. Stradner, and K. Crailsheim. Hormone-based control for multi-modular robotics. In P. Levi and S. Kernbach, editors, *Symbiotic Multi-Robot Organisms: Reliability, Adaptability, Evolution*. Springer-Verlag, February 2010.
- [56] T. Schmickl, H. Hamann, J. Stradner, R. Mayet, and K. Crailsheim. Complex taxis-behaviour in a novel bio-inspired robot controller. In H. Fellermann, M. Dörr, M. M. Hanczyc, L. L. Laursen, S. Maurer, D. Merkle, P.-A. Monnard, K. Støy, and S. Rasmussen, editors, *Proc. of the ALife XII Conference*, pages 648–655. MIT Press, 2010.
- [57] T. Schmickl, H. Hamann, and K. Crailsheim. Modelling a hormone-inspired controller for individual- and multi-modular robotic systems. *Mathematical and Computer Modelling of Dynamical Systems*, 17(3):221–242, 2011.
- [58] H.-P. Schwefel. *Evolution and Optimum Seeking*. Wiley & Sons, New York, USA, 1995. ISBN 0-471-57148-2.
- [59] W.-M. Shen, B. Salemi, and P. Will. Hormone-inspired adaptive communication and distributed control for CONRO self-reconfigurable robots. *IEEE Transactions on Robotics and Automation*, 18(5):700–712, Oct. 2002.
- [60] W.-M. Shen, P. Will, A. Galstyan, and C.-M. Chuong. Hormone-inspired self-organization and distributed control of robotic swarms. *Autonomous Robots*, 17:93–105, 2004.

- [61] W.-M. Shen, M. Krivokon, H. Chiu, J. Everist, M. Rubenstein, and J. Venkatesh. Multimode locomotion via SuperBot reconfigurable robots. *Autonomous Robots*, 20(2):165–177, 2006.
- [62] K. O. Stanley and R. Miikkulainen. Evolving neural networks through augmenting topologies. Technical Report AI2001-290, Department of Computer Sciences, The University of Texas at Austin, 2002. URL <http://nn.cs.utexas.edu/?stanley:ec02>.
- [63] K. O. Stanley and R. Miikkulainen. Competitive coevolution through evolutionary complexification. *Journal of Artificial Intelligence Research*, 21(1):63–100, January 2004.
- [64] K. O. Stanley, D. B. D’Ambrosio, and J. Gauci. A hypercube-based encoding for evolving large-scale neural networks. *Artificial Life*, 15(2):185–212, March 2009. doi: 10.1162/artl.2009.15.2.15202.
- [65] J. Stradner, H. Hamann, T. Schmickl, and K. Crailsheim. 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, 2009.
- [66] J. Stradner, H. Hamann, T. Schmickl, R. Thenius, and K. Crailsheim. Evolving a novel bio-inspired controller in reconfigurable robots. In G. Kampis, I. Karsai, and E. Szathmáry, editors, *Advances in Artificial Life, 10th European Conference, ECAL 2009*, volume 5777 of *Lecture Notes in Computer Science*, pages 132–139. Springer-Verlag, 2011.
- [67] SYMBRION. Project website, 2011. <http://www.symbion.eu>.
- [68] K. Takahashi, S. N. V. Arjunan, and M. Tomita. Space in systems biology of signaling pathways towards intracellular molecular crowding in silico. *FEBS Letters*, 579:17831788, 2005.
- [69] V. Trianni. *Evolutionary Swarm Robotics - Evolving Self-Organising Behaviours in Groups of Autonomous Robots*, volume 108 of *Studies in Computational Intelligence*. Springer-Verlag, Berlin, Germany, 2008.
- [70] P. A. Vargas, R. C. Moioli, L. N. de Castro, J. Timmis, M. Neal, and F. J. von Zuben. Artificial homeostatic system: a novel approach. In M. S. Capcarrère, A. A. Freitas, P. J. Bentley, C. G. Johnson, and J. Timmis, editors, *8th European Conference on Artificial Life (ECAL’05)*, volume 3630 of *LNAI*, pages 754–764. Springer-Verlag, 2005.
- [71] P. A. Vargas, R. C. Moioli, F. J. von Zuben, and P. Husbands. Homeostasis and evolution together dealing with novelties and managing disruptions. *International Journal of Intelligent Computing and Cybernetics*, 2(3), 2009.

- [72] S. Whiteson and P. Stone. Evolutionary function approximation for reinforcement learning. *Journal of Machine Learning Research*, 7:877–917, 2006.
- [73] B. Widrow and F. W. Smith. Pattern recognizing control systems. In J. T. Tou and R. H. Wilcox, editors, *Computer and Information Sciences*, pages 288–317. Clever Hume Press, Washington, DC, 1964.
- [74] N. Wiener. *Cybernetics: or Control and Communication in the Animal and the Machine*. MIT Press, Cambridge, MA, 1948.
- [75] L. Winkler and H. Wörn. Symbricator3D - A distributed simulation environment for modular robots. In M. Xie, Y. Xiong, C. Xiong, H. Liu, and Z. Hu, editors, *ICIRA*, volume 5928 of *LNCIS*, pages 1266–1277. Springer-Verlag, 2009.
- [76] D. H. Wolpert and W. G. Macready. No free lunch theorems for optimization. *IEEE Transactions on Evolutionary Computation*, 1(1):67–82, 1997.
- [77] X. Xin and M. Kaneda. Analysis of the energy based control for swinging up two pendulums. *IEEE Transactions on Automatic Control*, 50(5):679–684, 2005.
- [78] M. Yim, W. Shen, B. Salemi, D. Rus, M. Moll, H. Lipson, E. Klavins, and G. Chirikjian. Modular self-reconfigurable robot systems. *IEEE Robotics and Automation Magazine*, 14(1):43–52, 2007.