



HAL
open science

Flexible and Stable Pattern Generation by Evolving Constrained Plastic Neurocontrollers

Thierry Hoinville, C. Tapia, Patrick Henaff

► **To cite this version:**

Thierry Hoinville, C. Tapia, Patrick Henaff. Flexible and Stable Pattern Generation by Evolving Constrained Plastic Neurocontrollers. Adaptive Behavior, 2011, 19, pp.187-207. 10.1177/1059712311403631 . hal-00784897

HAL Id: hal-00784897

<https://hal.science/hal-00784897v1>

Submitted on 4 Feb 2013

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Abstract

In evolutionary robotics, plastic neural network models proved to be promising for evolving adaptive behaviors. In particular, neurocontrollers incorporating hebbian synapses have been shown to be useful for implementing conflicting sub-behaviors. Numerous interesting complex tasks assume such flexibility. However, those evolved controllers often exhibit behavioral instability, as simulation time is extended beyond the short limit used during evolution. In this paper, we propose constrained plastic models inspired by neural homeostasis phenomena, in order to evolve flexible and stable pattern generators for single-legged locomotion. Comparative results show that constrained controllers perform better than unconstrained ones in both terms of evolvability and behavioral stability. Functional analyses of the best evolved controller unveil the adaptivity, robustness and homeostasis arising from the statically constrained plasticity. Interestingly, homeostasis evolved implicitly without relying on any active homeostatic mechanisms and is implemented through hebbian plasticity, usually considered destabilizing.

Keywords

Homeostasis, plastic neural networks, adaptive synapses, legged locomotion, evolutionary robotics

I. Introduction

A key issue of evolutionary robotics is to seek convenient control building blocks for specific behaviors (Harvey, Husbands, Cliff, Thompson, & Jakobi, 1996; Jakobi & Quinn, 1998; Floreano & Urzelai, 2000). In particular, many efforts have been carried out in exploring new neural network models intended to be more evolvable, scalable and adaptive. Among them, Continuous Time Recurrent Neural Networks (CTRNNs) (Beer, 1995), GasNets (Husbands et al., 1998) and Plastic Neural Networks (PNNs) (Floreano & Mondada, 1996) received most attention. Comparing, hybridizing and refining these models, as initiated by McHale and Husbands (2004a, 2004b), turn out to be essential for better defining their areas of applicability (Bullock, 2006). Ultimately, it could allow us to predict which neural networks are suitable to evolve which behaviors. Here, we propose to study particular refinements applied to a CTRNN/PNN hybrid model.

The standard PNN model involves hebbian learning rules used to make each synapse adaptive. It was first introduced by Floreano and Mondada (1996) to explore evolution *of* plasticity, rather than evolution *and* plasticity. One of their later results (Urzelai & Floreano, 2001) was to show, through a Khepera performing a light-switching task, that evolved controllers embedding adaptive synapses can sequentially exhibit multiple possibly conflicting sub-behaviors. Such behavioral versatility constitutes a desired trait of autonomy and adaptivity. This contrasts with the monotonic and minimalistic behaviors usually obtained

¹Italian Institute of Technology, Genoa, Italy

²ETIS, University of Cergy-Pontoise, France

Corresponding author:

Dr. Thierry Hoinville Italian Institute of Technology Via Morego,
30 – 16163 Genoa Italy

Email: thierry.hoinville@iit.it

in prior works (Urzelai & Floreano, 2001; Doncieux, 2003). Subsequently, on evolving conflicting phototactic behaviors, Di Paolo (2003) obtained a similar result using another kind of adaptive-synapse controllers, based on spike-timing dependent plasticity (STDP).

However, PNNs face also several shortcomings. For example, Stanley, Bryant, and Miikkulainen (2003) reported that PNNs evolved slower than fixed-weight networks, for a dangerous food foraging task. Furthermore, Tuci and Quinn (2003) were unable to evolve PNNs to solve a learning task. This appears very surprising when one considers the many links observed in nature between hebbian plasticity and learning. Finally, for our concern in this paper, there are several evidences that PNN-controlled robots are inherently subject to some behavioral instability. First, none of the PNNs evolved by Urzelai and Floreano (2001) have retained permanently the desired light-switching behavior. That is, these controllers were unable to lead the Khepera to put the light on more than once, i.e. more than required during evolution. Second, PNNs evolved by McHale and Husbands (2004a, 2004b) for quadrupedal and bipedal locomotions tend to generate unstable gaits, as simulation time is extended beyond the limit used during evolution. We suggest that this last issue is an instance of the flexibility/stability dilemma. In other words, one can wonder how to evolve robot controllers that are plastic – thus providing some flexibility – and yet bearing behavioral stability?

In (Hoinville & Hénaff, 2004a, 2004b), we proposed to take inspiration from neurophysiological observations, which suggest that neurons counteract the destabilizing influence of hebbian plasticity using several homeostatic plasticity mechanisms (Turrigiano & Nelson, 2000). Practically, we have shown, for a single-legged locomotion task, that CTRNN/PNN hybrid models constrained by static rules inspired by homeostatic mechanisms are more evolvable and behaviorally stable than unconstrained ones. In this paper, we detail these results and extend them by performing functional analyses of the best evolved controller, so as to uncover the actual roles of the so-called homeostatic constraints we proposed.

This paper is organized as follows. The next section gives, first, the biological background that inspired our work and second, a state of the art of evolving homeostatic neurocontrollers for robots. The third section describes the neuronal models we compared in evolving single-legged locomotion controllers. Results are presented in the fourth section and discussed later in the fifth section. Then, we conclude in the last section by giving some perspectives.

2. Background

In this section, we draw the biological background that inspired our work and review how neuronal homeostasis has been addressed in evolutionary robotics.

2.1. Homeostasis in neurophysiology

Hebbian plasticity is known as a ubiquitous phenomenon involved in major cognitive and adaptive functions, like learning and memory. Moreover, it plays a key role in refining synaptic connections during brain development. Thus, by reinforcing correlation between simultaneously active neurons, hebbian plastic rules participate in the significant flexibility of nervous systems. However, raw modeling of this principle comes with serious instability issues, since reinforcing neuronal correlations increases likeliness of simultaneous activations and so on. Early modeling studies yielded many mathematical adjustments intended to cope with synaptic saturation impacting the basic Hebb's rule (Miller & MacKay, 1994). Some experimental observations supported more plausible advanced auto-stabilizing hebbian rules, like the BCM law (Bienenstock, Cooper, & Munro, 1982) or some forms of STDP (Kepecs, van Rossum, Song, & Tegnér, 2002).

A more general perspective has emerged around the notion of neuronal homeostasis. Homeostasis refers to all phenomena that contribute to preserve some entity's internal physiological state (e.g. temperature) from internal or external disturbances (e.g. weather) (Kitano, 2007). Applied to neurons, this notion can be defined as the capacity, of a neuron or a network, to maintain a characteristic activity (i.e. state).

In the following sections, we will focus on two homeostatic phenomena, namely excitability regulation and synaptic scaling, which have inspired the present work. These are believed to take a great part in counteracting the destabilizing effect of hebbian plasticity on neuronal activity (Turrigiano & Nelson, 2000). The first one operates on intrinsic properties of neurons for regulating their excitability. The second one occurs at synaptic level to regulate the total synaptic input strength of each neuron.

2.1.1. Excitability regulation. The excitability of a neuron can be defined as its probability to fire action potentials according to input signal levels it receives. Neuronal excitability can be illustrated by a sigmoidal curve that relates the total synaptic drive received in input to the mean firing rate obtained in output (Figure 1). According to Turrigiano and Nelson (2000), excitability regulation can be described as a process that shifts this curve to maintain a long term

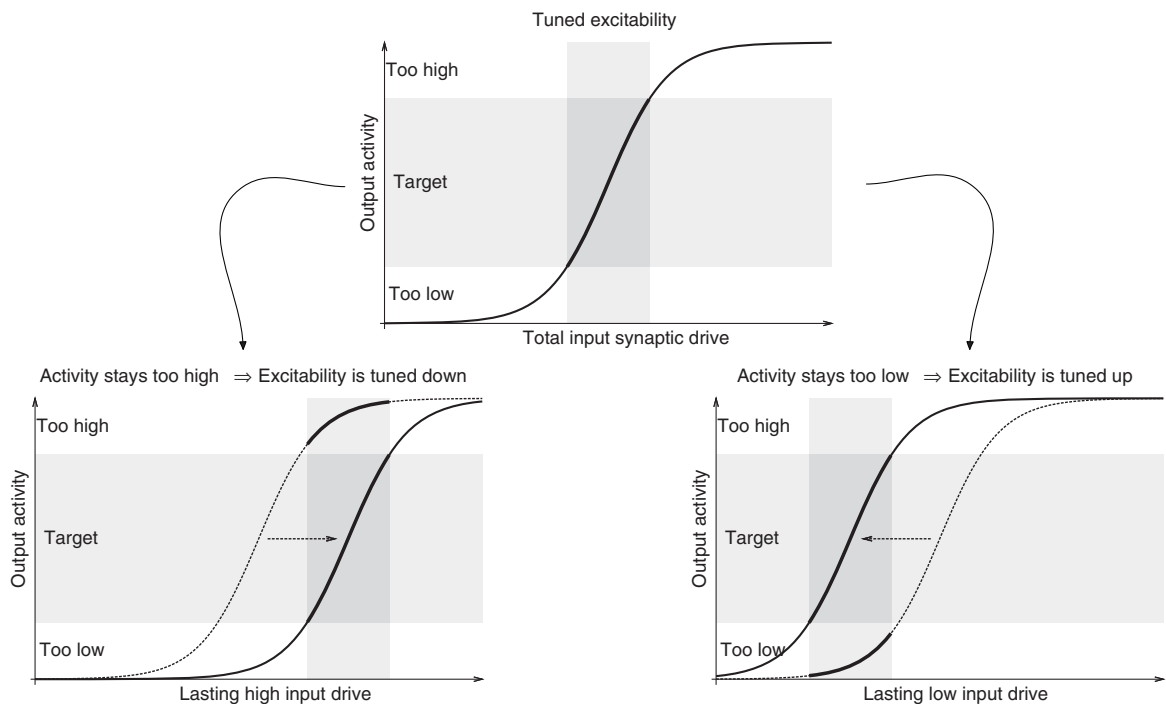


Figure 1. Homeostatic regulation of neuronal excitability. Typically, for a neuron, a sigmoid activation function can be used to relate the total synaptic drive received as input to the mean firing rate (i.e. activity) obtained as output. Depending on the range of synaptic drive received (vertical gray strips), only a part of the activation function is used (darkened curve parts). If this part corresponds to too high (left example) or too low (right example) neuronal activity, the activation function is shifted (i.e. excitability is tuned) in order to get back to a target activity level (horizontal gray strips).




Figure 2. Interaction between hebbian plasticity and synaptic scaling. One post-synaptic neuron is receiving excitatory inputs from two pre-synaptic neurons. When one of the latter fires enough for activating the post-synaptic neuron, the synapse involved is potentiated via hebbian plasticity, resulting in higher post-synaptic activity and so on. If this too high post-synaptic activity lasts too long, then synaptic scaling acts on all input synapses proportionally to their weights in order to get back to a normal activity. It leads to synaptic competition in that one potentiated synapse induces depression of the other.

“normal” activity. When input excitation induces too high firing rate for a long time, excitability is decreased (left example on Figure 1). Conversely, if input excitation is not sufficient, thus making the considered cell silent for a long while, then its excitability is increased (right example on Figure 1).

2.1.2. Synaptic scaling. This homeostatic phenomenon consists of scaling the global synaptic input gain

of a neuron, that counteracts prolonged perturbation of its activity (Figure 2). Thus, when activity remains too low for a long while, all excitatory inputs are potentiated and all inhibitory inputs are depressed. Opposite modifications are expressed when lasting over-activity occurs. Moreover, as reported by Turrigiano and Nelson (2000), synaptic scaling appears to be multiplicative, i.e. input synapses are potentiated or depressed in proportion to their strength. As such, synaptic

scaling has been shown to induce synaptic competition (Figure 2) and to overlay hebbian plasticity in a very complementary way (Turrigiano & Nelson, 2000).

2.2. Neuronal homeostasis in evolutionary robotics

2.2.1. Models and main results. Evolution of homeostatic neurocontrollers has been mostly performed by Di Paolo and his colleagues. The first model proposed in (Di Paolo, 2000) implements homeostasis through conditional activation of various hebbian plasticity rules. Some activation bounds are defined so as to divide neuronal activation in homeostatic and non-homeostatic ranges. Whenever any neuron fires out of the homeostatic bounds, genetically-specified hebbian laws are triggered on all its input synapses in order to recover a normal activity. This model has been used to evolve controllers for simulated Braitenberg-like agents equipped with two actuated wheels and two light sensors. The targeted task was to reach light sources in a 2D infinite space (this behavior will be referred in the following as “phototaxis”).

Then, this model has been slightly modified in further studies (two homeostatic ranges instead of one; other hebbian plasticity formulations). These ones still involved phototactic tasks, but following more complex scenarios designed to address some questions about cognition:

- modeling perseverative reaching in infants (Wood & Di Paolo, 2007);
- minimal dynamics for behavioral preference (Iizuka & Di Paolo, 2007);
- linkage between internal and behavioral stability (Iizuka & Di Paolo, 2008).

In another work aiming for simpler phototactic behavior, Di Paolo (2002a) proposed a model of “homeostatic oscillators” based on standard CTRNNs incorporating a fast mechanism of excitability regulation. This one consists of turning the bias of each neuron into a dynamic variable, which acts to keep constant its average firing rate to a given homeostatic value. Because of the fast speed of this plasticity, neurons tend to oscillate.

Di Paolo (2002b) also considered a more complex model based on spiking neurons, where each synapse is subject to both spike-timing dependent plasticity (STDP) and synaptic scaling. Basically, STDP drives synaptic weight changes according to the time difference between pre- and post-synaptic potentials. Among the compared models, one implements in addition to STDP a mechanism of synaptic scaling. Spike trains of each neuron are leaky integrated to estimate

the current firing rate which is then compared to a given reference firing rate. The difference is minimized by scaling up or down the weights of afferent synapses, depending on whether they are excitatory or inhibitory. Here again, the controllers based on these mechanisms were evolved to satisfy simple phototaxis.

All the above homeostatic models have shown improved evolvability compared to the control models, in term of both final scores and evolving speeds. The resulting neurocontrollers performed phototactic tasks more reliably and adaptively. For example, some proved to be robust to neural noise or randomization of initial state. Most of them were even able to adapt to strong sensorimotor disruptions never seen during evolution, like motor distortions, sensor swapping or removal (Di Paolo, 2000, 2002a, 2002b; Iizuka & Di Paolo, 2008). Besides, it appears that such homeostatic models provide some behavioral stability. In fact, except in (Di Paolo, 2000, 2002b), the obtained behaviors showed extended durability or readaptability (i.e. recovering from cancelling the sensorimotor perturbations).

In contrast, Williams (2005, 2006) reported poor evolvability results using another homeostatic network model. This one consists of a CTRNN model endowed with an excitability regulation and a synaptic scaling mechanism. Like in (Di Paolo, 2002a), excitability regulation is implemented by an adaptive bias. However, the dynamics involved here are slower and the plasticity just operates when neuronal activity goes outside a homeostatic range (i.e. activity is not constrained to a unique value as in Di Paolo’s model). In the same way, if a neuron fires outside this given homeostatic range, all its afferent synapses are multiplied down or up, depending on their sign and in way to return the neuron to a normal firing activity.

This model was assumed to be more evolvable than the classic CTRNNs because of the network-level properties it provides: increased sensitivity, better signal propagation and increased likelihood of oscillations (Williams & Noble, 2007). Nevertheless, evolution attempts made by Williams (2005, 2006) on a task of 1D catching falling objects (with or without shape discrimination) showed that homeostatic plastic CTRNNs were generally outperformed by non-plastic CTRNNs, although the formers evolved quicker and with more consistency. Williams (2005, 2006) showed however that such homeostatic plasticity, when used as developmental mechanism rather than lifetime plasticity, sometimes led to more evolvable controllers than non pre-developed ones.

2.2.2. Discussion. While enforcing the idea that neuronal homeostasis promotes evolution of adaptive

behaviors, the previous studies can be discussed in relation to their methods.

Firstly, in order to fit with short-duration simulations required for affordable fitness computation, all the modeled homeostatic mechanisms have been parameterized such as to act much faster than the inspiring biological processes. Although most authors ensured to significantly separate timescales of neuronal activation and homeostatic plasticity, we think that this approximation tends to denaturate the given mechanisms and lowers biological relevance of the results. For example, in (Di Paolo, 2002a) and (Di Paolo, 2002b), fast excitability regulation gives rise to rapid rhythmic neuronal activity. It hardly relates to nature considering the slow dynamics of the real mechanism, yet this function is useful here and evolutionary robotics in general may not be adapted to achieve biomimetic results.

Secondly, except in (Di Paolo, 2002b), homeostatic plasticity is envisaged as a solution to implement short-term behavioral adaptivity or robustness (which again is unrealistic given its slowness), rather than as a stabilizing counterpart to other plasticities, like it has been argued for the hebbian one (Turrigiano & Nelson, 2000).

Thirdly, the phototactic tasks may inherently favor the homeostatic models. In fact, even randomly generated homeostatic CTRNNs (i.e. before any evolution is applied) were shown in significant proportion to perform phototaxis when embodied in Braitenberg agents (Williams, 2004, 2006).

3. Methods

3.1. Models

All the models we have reviewed involve homeostatic plastic mechanisms that dynamically monitor neuronal activity and trigger some parameter changes to recover homeostasis when needed. In contrast, we propose to follow an approach based on static mathematical constraints intended to decrease the likelihood of neuronal activity being saturated or silent. These “homeostatic constraints” restrict the dynamical neural networks to parameter subspaces similar to those asymptotically reached using homeostatic plasticity. This keeps unchanged the dimension of the dynamics space (i.e. no new state variables) and discards the issue of slow timescale dynamics. We expect this would encourage evolutionary search to focus on richer dynamics (Mathayomchan & Beer, 2002) which also show proper behavioral stability, although in the presence of disturbing plasticity such as the hebbian one.

Furthermore, an open question is whether this approach could implicitly lead to evolve some homeostasis.

In particular, our present work investigates the impact of two specific homeostatic constraints on the evolvability, behavioral stability and the internal dynamics of CTRNN/PNN hybrid controllers. The first one mimics the outcome of excitability regulation. The second one is inspired from synaptic scaling.

3.1.1. CTRNN/PNN hybrid model. Firstly proposed in (McHale & Husbands, 2004a, 2004b) and (Hoinville & Hénaff, 2004a, 2004b), this model combines the leaky integrator neurons of CTRNNs with the hebbian synapses of PNNs.

Adapted from (Beer, 1995), neurons are modeled by

$$y_i^{t+\Delta t} = y_i^t + \frac{\Delta t}{\tau_{mi}} \left(-y_i^t + k_i \sum_{j=1}^N \pm_{ji} w_{ji}^t o_j^t + I_i^t \right) \quad (1)$$

$$i = 1, 2, \dots, N$$

where $y_i^{t+\Delta t}$ and y_i^t are the i th node’s states at respectively times $t + \Delta t$ and t ; Δt is the simulation time slice (set to 10 ms); τ_{mi} is the time constant of the i th node; $k_i = 1/(\sum_{j=1}^N |\pm_{ji}|)$ is a normalizing factor; \pm_{ji} is the sign of the connection from the j th to the i th nodes (1: excitatory, -1: inhibitory, 0: off); $w_{ji}^t \in [0, 1]$ is the corresponding connection strength which is subject to plastic changes; I_i^t is the input value in case of a sensory neuron, otherwise it is set to 0; o_j^t is the output of the j th node computed through a sigmoid activation function, such as

$$o_j^t = \frac{1}{1 + \exp(-g_j(y_j^t + \theta_j))} \quad (2)$$

with g_j and θ_j are respectively the j th node’s gain and bias.

According to the PNN model (Floreano & Mondada, 1996), connection weights are updated through time as follows:

$$w_{ji}^{t+\Delta t} = w_{ji}^t + \frac{\Delta t}{\tau_{sji}} \Delta w_{ji}^t \quad (3)$$

where τ_{sji} is the plasticity time constant (the equivalent learning rate is $\eta_{ji} = \Delta t/\tau_{sji}$) of the connection from the j th to the i th nodes and Δw_{ji}^t is one of the following four hebbian rules:

1. Plain Hebb rule

$$\Delta w_{ji}^t = (1 - w_{ji}^t) o_j^t o_i^t \quad (4)$$

2. Post-synaptic rule

$$\Delta w_{ji}^t = w_{ji}^t (-1 + o_j^t) o_i^t + (1 - w_{ji}^t) o_j^t o_i^t \quad (5)$$

3. Pre-synaptic rule

$$\Delta w_{ji}^t = w_{ji}^t o_j^t (-1 + o_i^t) + (1 - w_{ji}^t) o_j^t o_i^t \quad (6)$$

4. Covariance rule

$$\Delta w_{ji}^t = \begin{cases} (1 - w_{ji}^t) \mathcal{F}(o_j^t, o_i^t) & \text{si } \mathcal{F}(o_j^t, o_i^t) > 0 \\ w_{ji}^t \mathcal{F}(o_j^t, o_i^t) & \text{otherwise} \end{cases} \quad (7)$$

where the function $\mathcal{F}(o_j^t, o_i^t) = \tanh[4(1 - |o_j^t - o_i^t|) - 2]$ measures the similarity between pre- and post-synaptic activities.

3.1.2. Excitability regulation by center-crossing constraint.

Homeostatic regulation of neuronal excitability can be achieved by providing the sigmoid activation function with an adaptive bias. Such a solution has been carried out in (Di Paolo, 2002a; Williams & Noble, 2007). Moreover, there is evidence that this mechanism tends to maintain CTRNNs close to the so-called ‘‘center-crossing position’’ of the parameter space (Williams & Noble, 2007; Williams, 2006).

Center-crossing CTRNNs are defined as satisfying a special condition on neuronal bias given by

$$\theta_i^* = \frac{-\sum_{j=1}^N \pm_{ji} w_{ji}^t}{2} \quad (8)$$

which insures that the operating range of each neuron is centered on the most sensitive region of its activation function (Mathayomchan & Beer, 2002).

Here, we propose to use this condition to constrain the neurons, rather than possibly reach homeostatic states through dynamic implementation of excitability regulation. However, setting θ_i to θ_i^* is quite restrictive and reduces the number of parameters. Therefore, to keep the parameterization unchanged between our models and let evolution seek the nearby center-crossing condition, we chose the more flexible solution of correcting the activation function asymmetry:

$$y_i^{t+\Delta t} = y_i^t + \frac{\Delta t}{\tau_{mi}} \left(-y_i^t + k_i \sum_{j=1}^N \pm_{ji} w_{ji}^t (2o_j^t - 1) + I_i^t \right) \quad (9)$$

Indeed, for symmetric odd activation functions such as $(2o_j^t - 1)$, the center-crossing condition is verified for $\theta_i^{**} = 0$.¹ So, in this model, by both keeping evolving θ_i and symmetrizing the activation function (yet just for neuronal summation, not for output), we center the evolutionary search around the center-crossing condition and keep our models homogeneous.

3.1.3. Synaptic scaling by normalization constraint.

As we have seen before, synaptic scaling regulates global input gain of neurons to maintain a typical activity. Besides, it can lead to synaptic competition that properly complements hebbian plasticity.

Normalization constraints have been proposed in the past (Miller & MacKay, 1994) for modeling synaptic competition that takes place notably during ocular dominance column development.² The principle is to keep either $\sum_{j=1}^N \pm_{ji} w_{ji}^t$ or $\sum_{j=1}^N (w_{ji}^t)$ constant, the latter giving usually more uniform synaptic distributions (Miller & MacKay, 1994).

We propose to use $\sum_{j=1}^N (w_{ji}^t) = 1$. Indeed, we assume that constraining synaptic input gains to unity would balance information propagation throughout neurons and so would promote maintaining meaningful neuronal activity, i.e. not saturating, neither silent.

To be consistent with synaptic scaling observations, we implemented the normalization constraint in a multiplicative way, as follow:

$$w_{ji}^{t+\Delta t} = \frac{w_{ji}^t + \frac{\Delta t}{\tau_{sj}} \Delta w_{ji}^t}{\sqrt{\sum_{k=1}^N \left(w_{ki}^t + \frac{t}{\tau_{ski}} w_{ki}^t \right)^2}} \quad (10)$$

and

$$y_i^{t+\Delta t} = y_i^t + \frac{\Delta t}{\tau_{mi}} \left(-y_i^t + \sqrt{k_i} \sum_{j=1}^N \pm_{ji} w_{ji}^t o_j^t + I_i^t \right) \quad (11)$$

3.1.4. Summary.

It is important to note that, contrary to the usual approach, we do not introduce any new dynamics for enforcing homeostasis. That is, the neuronal activity is not monitored and there is no mechanism that dynamically corrects any parameter. In fact, homeostasis is not ensured to be maintained in the short- or long-term. However, our hypothesis is that the chosen constraints would make homeostasis more likely and lead to evolve plastic controllers implementing stable behaviors.

To evaluate this hypothesis, we compared four neuronal models instantiated from the two homeostatic constraints defined above:

- CTRL, the plain CTRNN/PNN hybrid model (our control model);
- CC, the CTRNN/PNN model with center-crossing neurons;
- NS, the CTRNN/PNN model with normalized adaptive synapses;
- CCNS, the CTRNN/PNN model with both center-crossing neurons and normalized adaptive synapses.

3.2. Application

We compared the performances of these four models as substrates for evolving controllers for a single-legged locomotion task. Such a test bed is preferred to the

phototactic tasks usually found in related works for two main assumptions.

Firstly, as we pointed out in the background section, phototactic tasks might inherently suit to homeostatic networks, as the success of random homeostatic CTRNNs in (Williams, 2004, 2006) tends to show. Although adaptation to unseen perturbation like sensory swapping (Di Paolo, 2000) requires non trivial dynamics, phototaxis of Braitenberg-like agents can rely at low level on basic sensory-motor reactions. On the contrary, even involving a single leg, locomotion requires generating rhythmic pattern, possibly from scratch, for coordinating several degrees of freedom. Possible control strategies are multiple since they rely on adapting any combination of frequencies, amplitudes and phases.

Secondly, considering homeostasis, maintaining a “normal” oscillatory activity underlying some locomotor behavior appears challenging. Defining simple homeostatic activity ranges, as usually proposed, would not be enough to keep stable a specific neuronal activity pattern. To be effective, active homeostatic mechanisms would thus require to monitor several oscillatory signal properties such as phases and frequencies. Here, our approach based on homeostatic parameter constraints avoids this issue.

3.2.1. Single-legged robot. The simulated robot² we used consists in a cuboid base supporting a leg articulated about three degrees of freedom (Figure 3). The hip and knee are universal and hinge joints respectively, allowing protraction/retraction, elevation/depression and extension/flexion of the leg.

In addition, a linear joint between the base and the ground constrains the robot to move along a horizontal straight line (Y -axis, 12 cm high). This rail generates viscous friction that brakes the robot sliding. The friction force $F_{fr}(t)$ is proportional to the base’s instantaneous speed $V(t)$, such as

$$F_{fr}(t) = -k_{fr}V(t) \quad (12)$$

where k_{fr} is the viscous friction coefficient.

3.2.2. Fitness. The task of the robot is to respect an instantaneous desired velocity $V_d(t)$. In addition, the robot must cope with an external perturbation that consists in sudden increase of the friction coefficient k_{fr} .

In order to evolve controllers satisfying this behavior, we defined a fitness based on three evaluation scenarios. As shown in Table 1, each scenario is defined by time-varying profiles of the desired speed $V_d(t)$ and the friction coefficient $k_{fr}(t)$:

- the scenario A aims to follow a triangular desired speed profile;

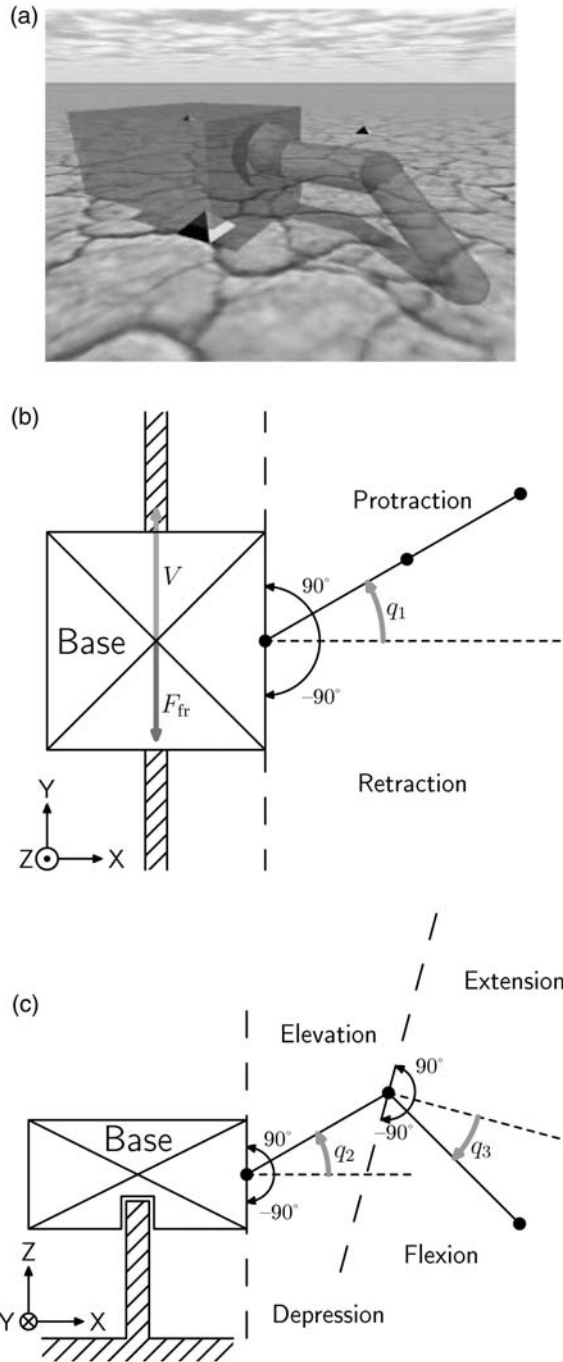
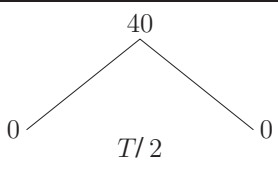
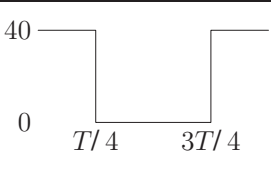

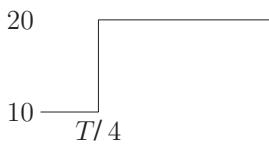


Figure 3. The single-legged robot model. The base is a box of $20 \times 20 \times 10$ cm and 3 kg. The leg segments are cylinders of $15 \times \varnothing 4$ cm and 0.5 kg. In the initial configuration (i.e. $q_1, q_2, q_3 = 0$), the thigh is parallel to the ground and the tibia is bent to 45° toward the ground. From this configuration, all rotation ranges are $\pm 90^\circ$. (a) Simulation (b) Top view (c) Back view.

Table 1. The three evaluation scenarios ($T = 10$ s)

	Scenario A	Scenario B	Scenario C
$V_d(t)$ [cm/s]			
$k_{fr}(t)$ [kg/s]	10	10	

- the scenario B rewards the capacity to inhibit the locomotion behavior;
- the scenario C promotes adapting to the perturbation.

Each scenario involves simulating $T = 10$ seconds of real time (i.e. 1000 timesteps), during which the mean error E between the desired velocity $V_d(t)$ and the instantaneous mean velocity $V_m(t)$ is computed according to

$$E = \frac{1}{T} \int_0^T |V_d(t) - V_m(t)| dt \quad (13)$$

where $V_m(t)$ is measured through applying a second order low-pass filter on the raw velocity $V(t)$. This filtering discards the high amplitude variations of $V(t)$ occurring at each stride of the robot.

From the mean errors E_A , E_B , E_C obtained in the three scenarios, the fitness of a controller, which should be minimized, is given by:

$$fitness = \sqrt{E_A^2 + E_B^2 + E_C^2} \quad (14)$$

This quadratic combination encourages minimizing error in all the three scenarios by reducing the local optima corresponding to solving one or two scenarios only.

This fitness function has been built empirically by several evolution trials. Gathering these three scenarios allows to reward controllable and generalizing closed-loop locomotor behaviors. Note that, since each scenario lasts for a short time of ten seconds, the fitness function does not reward for long-term behavioral stability.

3.3. Controllers architecture and genetic encoding

For all the four models, we used the same control architecture of eight neurons (Figure 4). As we will describe further, the network topology was encoded in the genotype of each controller. We only kept fixed the connections with sensors and actuators. The neurons no. 1 and 2 are sensorineurons. The first is fed by the instantaneous error $e(t) = V_d(t) - V_m(t)$ between the desired velocity and the current average velocity. The second neuron receives boolean value from a ground contact sensor at the tip of the leg. Activations of the neurons no. 6, 7 and 8 give, after linear scaling, the angles of the three articulations³. The three remaining neurons are hidden interneurons.

As our controller models share the same parameterization and architecture, we used a common genetic encoding scheme (Figure 5). Each genotype is composed of eight neuron blocks. Each neuron block i includes the intrinsic properties (i.e. τ_{mi} , g_i and θ_i) of the coded neuron and a list of eight synapse blocks. These are coding for the neuron inputs (including the connection from itself). Each synapse block j of the neuron block i codes: whether the connection exists or not (thereby it defines the network topology), the synapse sign \pm_{ji} , one of the four hebbian rules and the plasticity time constant τ_{sji} .

In total, each genotype consists in 280 genes. Each numerical gene can take one value among five (Table 2). Hence, there are 3.74×10^{138} possible genotypes in the search space.

3.4. Genetic algorithm

We evolved the controllers using a classical generational and elitist genetic algorithm (Goldberg, 1989). That is, populations were initially randomly generated and

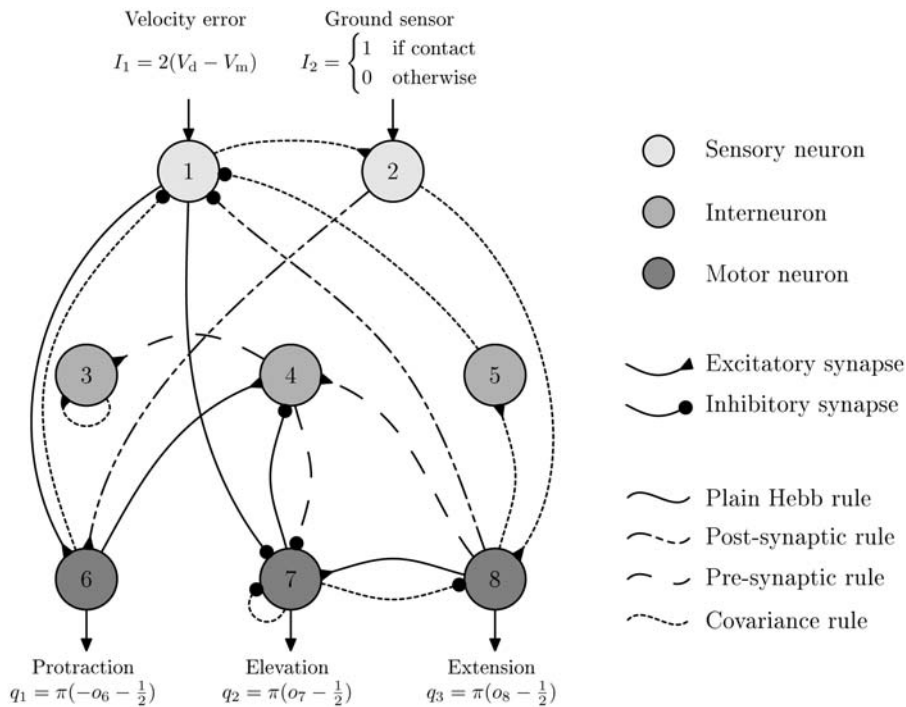


Figure 4. Evolved controller phenotype (corresponding to the best CCNS controller later analyzed).

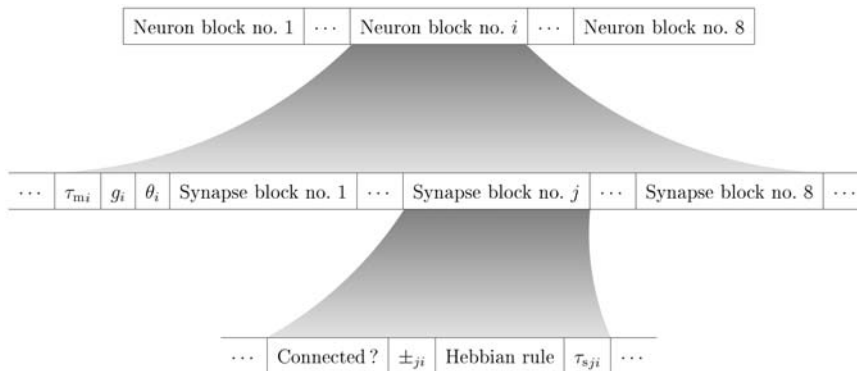


Figure 5. Genetic encoding scheme.

entirely renewed at each generation (except the best performing individual). To avoid premature convergence, we used stochastic universal sampling (Baker, 1987) and linear ranking selection with selective pressure of 1.1 (Goldberg, 1989). New individuals were generated using two genetic operators: simple allelic mutation ($P_{mut} = 0.001$) and uniform crossover ($P_{cross} = 0.6$).

4. Results

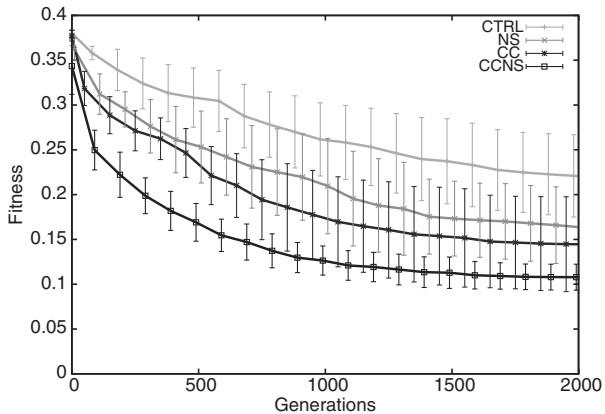
For each of the four controller models, we performed ten independent evolutionary runs, initialized with distinct populations of 200 random individuals. Each evolution was conducted for 2000 generations.

4.1. Evolvability

Figure 6 shows, for each model, the fitness profile of the best controllers averaged across all evolutionary runs. The non-homeostatic model (CTRL) is less evolvable than all the homeostatic models (CC, NS, CCNS) both in terms of final result and convergence speed (p -values $< .01$). Final performances of the CC and NS models are not significantly different (p -value $> .1$), meaning that one homeostatic constraint does not outperform the other. In contrast, the CCNS model combining both constraints appears to be the most evolvable especially regarding results consistency

Table 2. Allele sets of numerical genes

Genes	Alleles				
τ_m	0.02	0.165	0.31	0.455	0.6
g_i	2.46	3.53	5.34	9.43	31.26
θ_i	-0.2	-0.1	0.0	0.1	0.2
τ_{s_j}	0.2	2.65	5.1	7.55	10.0

**Figure 6.** Average best fitness for each neuronal model.

(indicated by the lower standard deviation) and convergence speed. Although, considering average final performances, the difference between the CC and CCNS model is not significant enough (p -value = .06).

4.2. Behavioral stability and robustness

In order to verify whether the homeostatic constraints promote behavioral stability or not, we evaluated the fitness of every best controller in time-extended simulation lasting ten times the evolutionary one (i.e. $T=100$ seconds). Figure 7 shows the obtained global fitness and errors specific to each scenario.

Leading to the lowest error values, the scenario A appears to be the least challenging. For all models, prolonged simulations drop the global fitness. This is especially marked for the CC model (whose fitness got worse by 72%) and less important for the CTRL, NS and CCNS models (resp. 43%, 33% and 20% worse). In term of absolute performance, the superiority of the CCNS is confirmed, whereas the CTRL model displays the worst results. Again, differences between the CC and NS models are not significant, although we notice a slight but systematic inversion in their relative fitness whatever the scenario, compared to evolutionary fitness.

To focus on robustness to the friction perturbation, we further evaluated the behavioral stability of the best

controllers using a custom test based on the scenario C. In that test, controllers were given a constant desired speed $V_d=30$ cm/s. A controller would succeed in the test if the robot reached the distance of 150 meters in less than 1000 seconds. Notice that, for a perfect controller continually respecting the desired velocity, this distance should be covered in 500 seconds. Compared to the previous fitness-based test, this custom test is more demanding on behavioral stability, since it involves longer simulated time.

We performed this test in perturbed condition and control condition. In the latter, we held constant the friction coefficient at $k_{fr}=10$ kg/s, whereas with perturbation, we increased it twofold such as $k_{fr}=20$ kg/s at the instant $t=250$ s.

Table 3 shows the success rate and average time-to-goal of each controller type in both environments. These results confirm and refine the previous ones. Showing almost optimal performances in both environments, CCNS controllers are more stable and robust than the others. The CTRL model gives the worst results, especially poor in the perturbed environment. CC and NS controllers display similar average success and good time-to-goal as long as the perturbation does not occur. However, in the perturbed environment, the success rate of the CC model drops critically, whereas NS controllers performances are only slightly affected, thus showing a great robustness (comparable to the CCNS one).

4.3. Behavioral analysis

In order to understand why homeostatic constraints led to improved performances, we analyzed both the behavior and internal dynamics (in the next section) characterizing the best obtained CCNS controller (whose phenotype is depicted on Figure 4). We first verified the capacity of this controller to follow entirely new desired velocity profiles, with and without perturbations. For instance, Figure 9a shows that the controller is able to respect a sinusoidal command while being perturbed. In Figure 9b, the square signal example clearly demonstrates the three basic capabilities of “speed control”, “stop-and-go” and “perturbation compensation”, which correspond to the three evolutionary scenarios.

To achieve such behaviors, the controller strategy is the following. For propulsion, the three leg’s degrees of freedom oscillate in phase at a constant frequency about 2 Hz. Amplitudes are small compared to the available ranges. Very simply, the locomotion cycle splits in two alternating phases: a stance phase (protraction, elevation and flexion) and a swing phase (retraction, depression and extension).

For braking, the controller uses a passive method. Thus, when the actual velocity is high above the desired




Figure 7. Mean performances (+ std dev.) of the best controllers in long-lasting scenarios ($T = 100$ s). For every model, each best controller of each evolutionary run was evaluated ten times with different random initial synaptic weights.

Table 3. Success rate of the best controllers in reaching 150 m distance at desired speed $V_d = 30$ cm/s, without and with friction perturbation (i.e. $k_{fr} = 10 \rightarrow 20$ kg/s at $t = 250$ s). Average times-to-goal of the successful controllers are also reported. For every model, each best controller of each evolutionary run was evaluated ten times with different random initial synaptic weights

Models	Without perturbation		With perturbation	
	Success [%]	Time-to-goal [s]	Success [%]	Time-to-goal [s]
CCNS	90	537	80	586
NS	72	597	70	669
CC	64	576	37	675
CTRL	55	792	20	849

velocity, the controller triggers a reflex similar to a high amplitude swing phase. This lifts the leg off the ground and let the friction force brakes the robot advance. Quickly after the reflex occurred, the leg stabilizes off the ground in a stationary configuration ($q_1 = -4^\circ$, $q_2 = 21^\circ$, $q_3 = -76^\circ$).

For compensating the friction perturbation, as well as adjusting the robot velocity to the desired value, the best CCNS controller uses a common strategy. The latter consists in shifting and modulating the amplitude of both protraction and extension oscillations. As shown in Figure 8, this results in modulating the stride length. Then, because the gait frequency remains constant, it ends up by either changing the overall robot speed or compensating the varying friction force.

4.4. Functional analysis

Functional analyses were conducted on the same CCNS controller, in order to understand how the plastic and homeostatic mechanisms it integrates work in

concert for giving rise to the efficient behavior previously described.

4.4.1. Neuronal and synaptic dynamics. The first question we addressed was how the controller’s internal dynamics correlate to the exhibited behavior. Figure 10 shows timeplots of the neurons activities and synaptic weights obtained for the square velocity command of Figure 9(b). None of the eight neurons is saturating or silent. Synaptic weights spread uniformly over the whole value range. Propulsion and stopping phases are characterized by two clearly different synaptic configurations. During propulsive behavior, both neuronal and synaptic dynamics are oscillatory, at a unique frequency matching the gait cycle. In contrast, during stops, both neuronal and synaptic configurations are mostly stationary. In addition, whether the friction perturbation occurs or not during propulsive phases also influences controller’s dynamics, yet to a small extent.




Figure 8. Trajectories of leg's knee and tip, with and without the friction perturbation. (a) Side view (b) Top view.

4.4.2. Synaptic plasticity and homeostatic constraints. In order to uncover the specific roles of synaptic plasticity and each homeostatic constraint, we disrupted each mechanism independently. Here again, we used the square velocity example.

Figure 11(a) shows the behavior arising from a random synaptic configuration not ever modified by any plastic rules (yet respecting both the NS constraint and the sign \pm_{ji} in the genotype). Whatever the random configuration generated, the controller was not able to properly control the robot speed. However, it likely led to oscillatory regimes moving the robot. Moreover, oscillation frequencies obtained were often similar to the normal controller's one.

We then investigated freezing the synaptic plasticity during both propulsive and stopping phases (Figure 11(b) and 11(c)). In the former case, the CCNS controller is still capable of respecting a high velocity, yet compensating a bit worse for the friction perturbation. But, interestingly, the controller is then unable to inhibit entirely the gait when the command becomes null. In the case of freezing the plasticity during a stopping phase, gait inhibiting is still performed well, whereas the controller is not more able to reach a high speed command.

Using the same method, we tested relaxing each homeostatic constraint in different contexts. For both constraint relaxations, consequences did not depend on the moment (i.e. on the controller state) we elicited them. Thus, here we only report results obtained for initial suppression of each constraint.

As shown in Figure 12, relaxing one constraint or the other induces totally distinct effects. On one hand, removing synaptic normalization (NS) still results in a working controller yet it has poor performance. In particular, the controller is unable to stop the robot when asked. In addition, as one could expect, this saturates several synapses. In contrast, as for neuronal activity, no saturation or silencing is induced. On the other hand, removing the center-crossing constraint (CC), by switching from the mixed equation (9) and equation (11) to equation (11), results in a controller that does not work at all. Logically, most neurons become either saturated or silent, paralysing the whole robot.

4.4.3. Plasticity dynamics. We further investigated how the plasticity dynamics impacts the emerging behavior, by achieving two more experiments on the same best CCNS controller.

First, we increased tenfold every plastic time constant $\tau_{s_{ji}}$ to slow down the plasticity dynamics. Figure 13(a) shows that this modification does not significantly affect the asymptotic behaviors. In particular, notice that the frequency of oscillatory regimes remains unchanged. However, behavioral transitions take much more time to finish and are easier to observe. Interestingly, this confirms that compensating the friction perturbation relates more to small parameter adjustments, than to an actual change in the dynamics.

Second, we tested whether individual differences in the dynamics of adaptive synapses were relevant to proper operation of the controller. Thus, we set all the plastic time constants to their mean value $\bar{\tau}_s = 1.56$ seconds. Results (Figure 13(b)) show that the controller is still able to generate a performing walking gait. However, some instability arises in the presence of the friction perturbation. Furthermore, the controller is unable to inhibit the locomotion cycle when subjected to a null command.




Figure 9. Velocity control and joint angles of the best CCNS controller with and without friction perturbation: (a) Sinusoidal velocity command of period 10 seconds; (b) Square velocity command of period 24 seconds. Instantaneous mean velocity V_m is oscillating (yet damped by the low-pass filter) according to every robot stride. An initial transitory phase lasting three seconds is noticeable. Gray stripes denote joint angle amplitudes characterizing the locomotion at $V_d = 30$ cm/s without friction perturbation (i.e. $k_{fr} = 10$ kg/s). Applying and removing the friction perturbation (highlighted by the vertical dash lines) induces undershoots and overshoots of the robot velocity respectively, which are quickly corrected by the controller.

4.4.4. Robustness. We ended our functional analysis of the best CCNS controller by studying its robustness to internal perturbation. While interacting with the environment, we disrupted the whole synaptic configuration by setting every synaptic weight to a new

random value. As shown in Figure 14, the controller proved to be very robust whatever the context in which we did the randomization. In every instance, the performance, activity and synaptic steady states are quickly recovered (in two to three seconds).

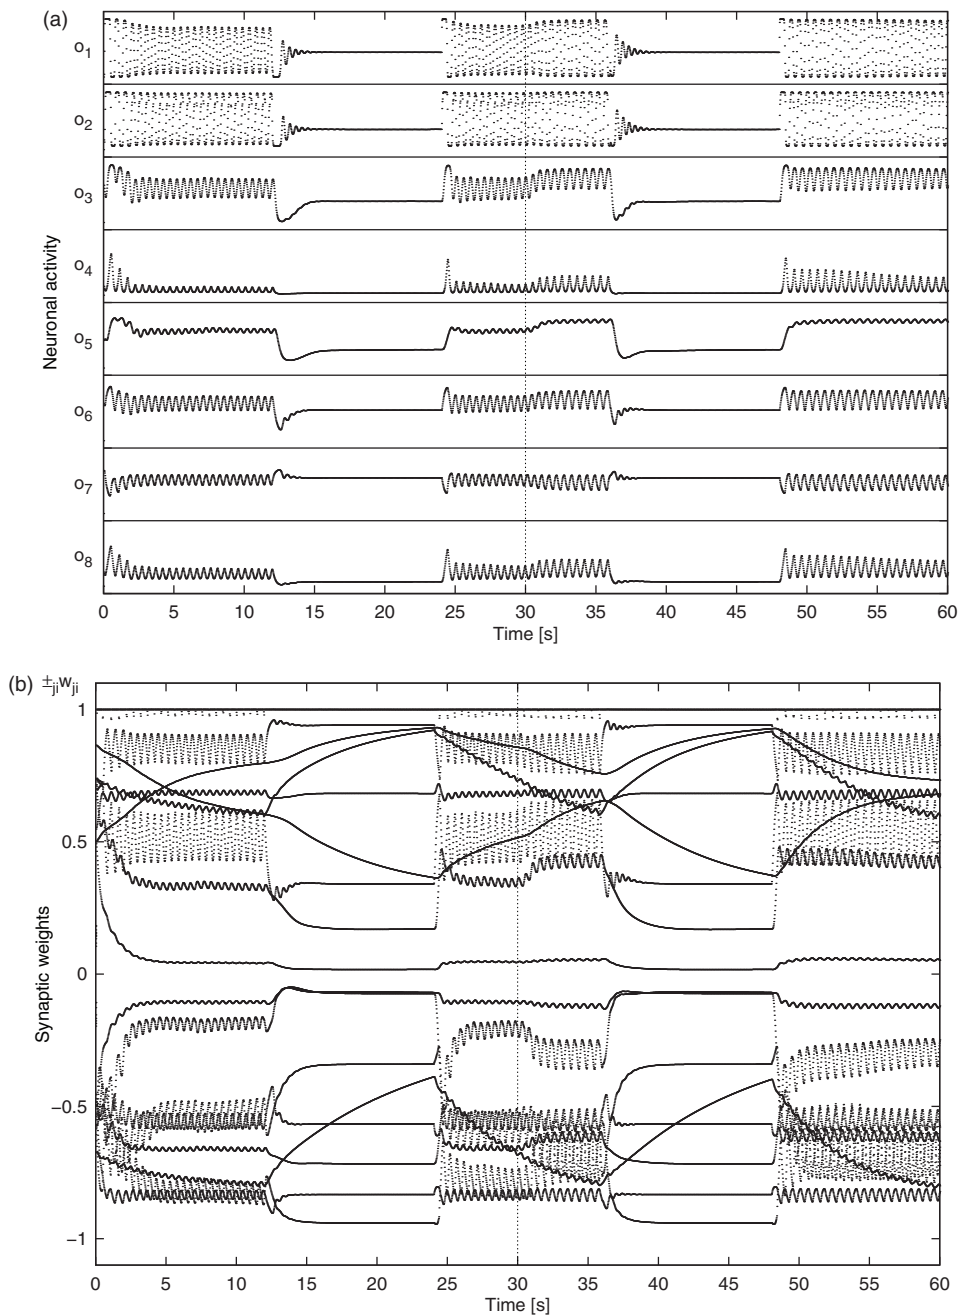


Figure 10. Dynamics of neuronal activity (a) and synaptic weights (b) of the best CCNS controller in the case of the square velocity example (Figure 9(b)). As two neurons (no. 2 and 5) each receives only one synaptic input, the normalization constraint keeps their weights fixed to 1.

5. Discussion

5.1. Hebbian plasticity underlies behavioral adaptivity

From the plasticity blocking tests performed on the best CCNS controller (Figure 11), we can first infer that adaptive synapses are not taking part in oscillatory

pattern generation. Actually, the contrary would have been possible given the relatively fast dynamics of the learning rules (yet slower than the neuronal activation dynamics). Pattern generation arises instead from the particular network topology and tuning of neuronal time constants, as confirmed by the fact that random fixed synaptic configurations still lead to effective propulsive behavior (Figure 11(a)). Furthermore, changes in plasticity speed (Figure 13) do not stop

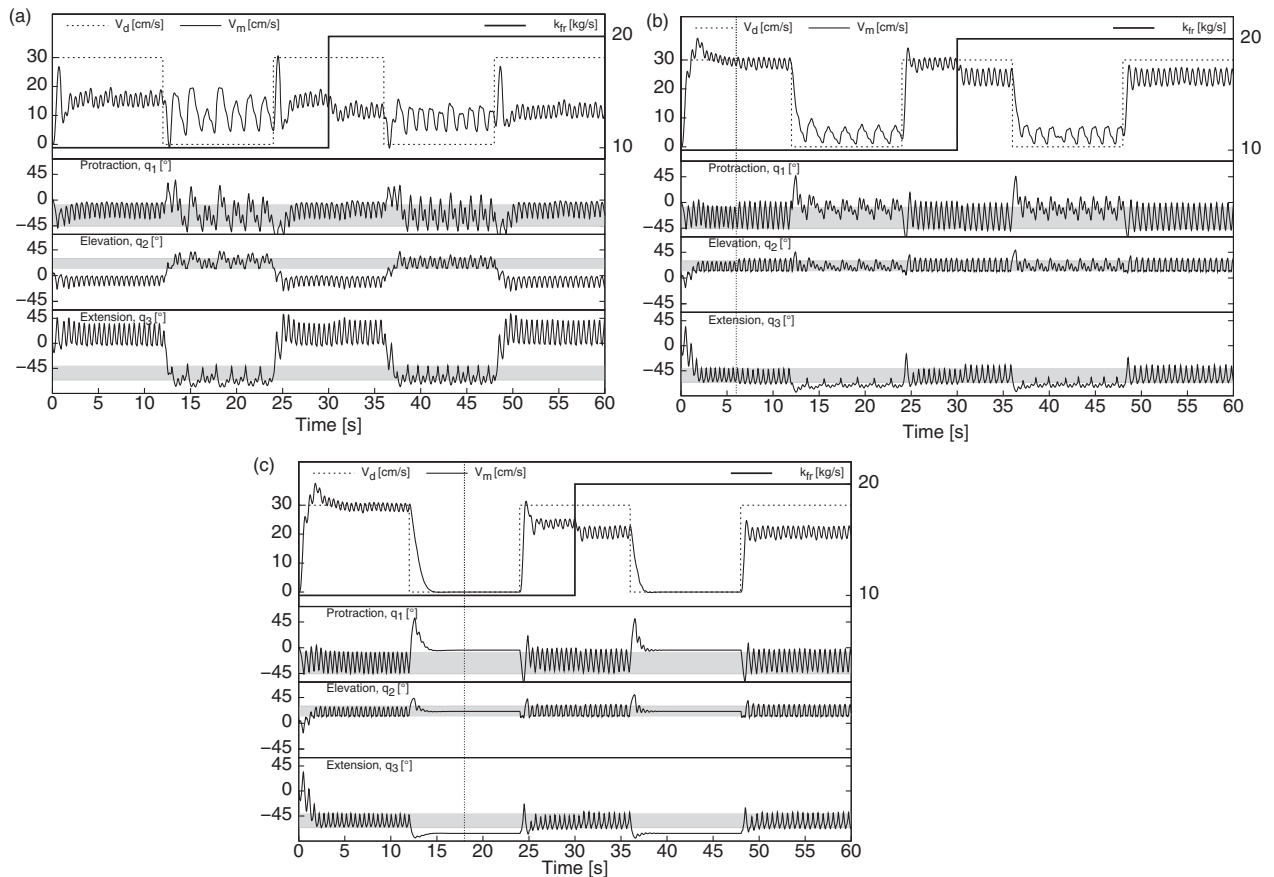


Figure 11. Effects of synaptic plasticity suppression: (a) Keeping the random initial synaptic configuration (yet normalized) makes the controller unable to control the robot speed, but a locomotion behavior is still exhibited. (b) Synaptic plasticity frozen during propulsion phase (from $t=6$ s): The controller is then unable to inhibit the robot locomotion. (c) Synaptic plasticity frozen during stop phase (from $t=18$ s): The controller is then unable to fully respect a high speed command. (Gray stripes corresponding to the intact controller behavior are reported here and in subsequent figures for comparison).

gait generation, neither do they modify oscillation frequency or phase.

The plasticity blocking tests however clearly show the role of synaptic plasticity in on-line behavior adaptation. The controller displays two distinct dynamics according to the required velocity: one for high speed propulsion and one for movement inhibition. Freezing all the synaptic weights in one dynamics prevents switching afterwards to the other (Figure 11). Besides, this relation between behavior transitions and internal plasticity becomes obvious when slowing down all adaptive synapses (Figure 13(a)).

On the other hand, adaptivity is less apparent when the controller compensates for the friction perturbation. In fact, small variations in synaptic weights are noticeable (Figure 10(b)) and induce only some changes in leg movement amplitude (Figure 8), whereas frequency and phase remain constant. Moreover, blocking the plasticity while being in the

propulsive regime without perturbation does not impair the later controller performance significantly, when the friction coefficient is to be doubled (Figure 11(b)). Therefore, plasticity does not seem involved in compensating such a perturbation. One possible explanation might be that the friction perturbation is not that demanding. Another explanation, which may add to the previous one, lies in that, in our setup, controllers could hardly detect this perturbation. Indeed, both friction perturbation and some increase in required velocity lead to similar rise in the control error perceived by controllers. In this way, a null velocity command makes a greater perturbation for controllers to deal with.

To summarize, the plasticity due to adaptive synapses is involved in adapting the control to big changes in required velocity, rather than in compensating the probably trivial friction perturbation. Consistent with previous work (Urzelai & Floreano, 2001), this result

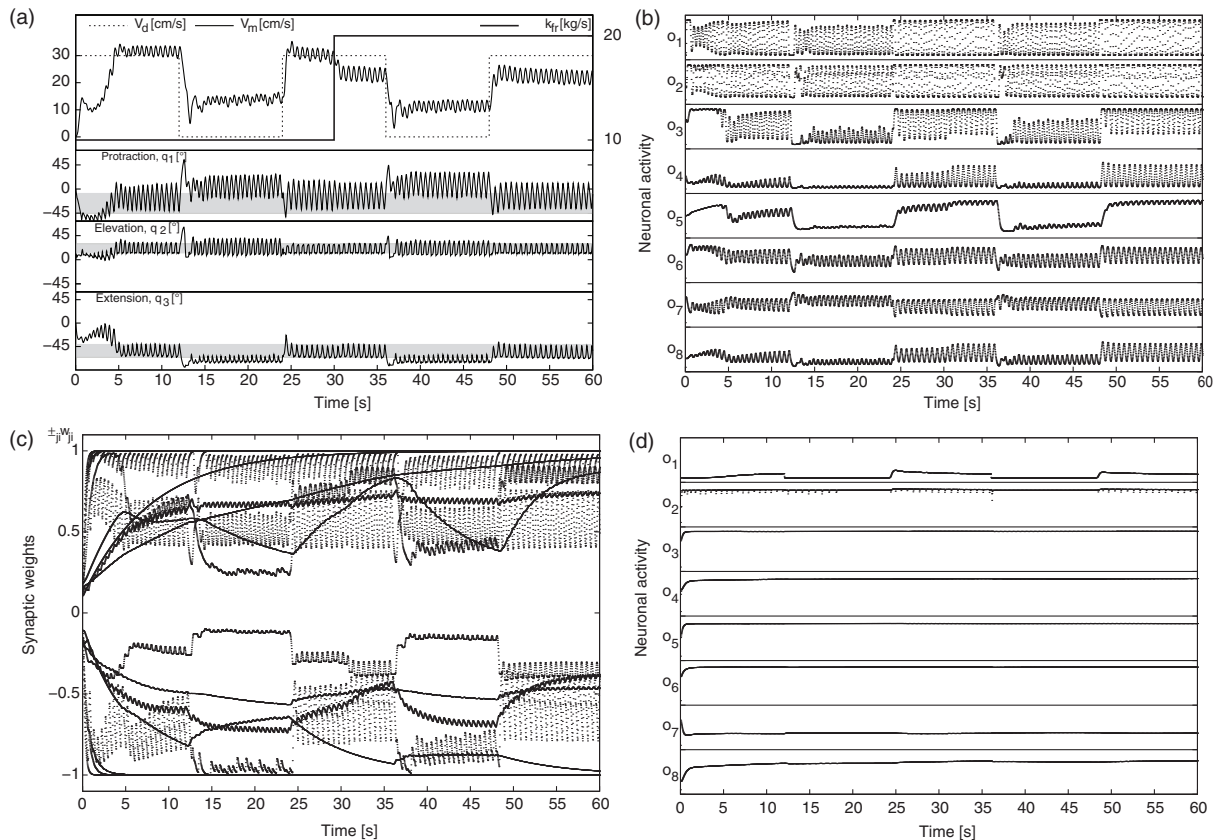


Figure 12. Effects of homeostatic constraint suppression: (a) Synaptic normalization constraint relaxation: The initial transitory regime is lengthened. The controller is unable to stop the robot. A significant part of synaptic weights is saturating (c), whereas neuronal activity remains oscillatory without saturating (b). (d) Center-crossing constraint relaxation: Almost all neurons display flat saturated activity, resulting in a non functional controller.

supports that the adaptive synapses model allows evolving controllers implementing conflicting sub-behaviors.

5.2. Synaptic normalization supports multistability

As pointed out in the introduction, the adaptivity of controllers evolved with adaptive synapses (namely the PNN model) has been offset by some behavioral instability before (McHale & Husbands, 2004a, 2004b; Urzelai & Floreano, 2001). This drawback is confirmed by our stability tests (Table 3) performed on the CTRL controllers (i.e. unconstrained CTRNN/PNN hybrid model), which display poor results especially in perturbed environment. On the other hand, evolving controllers with normalized synapses (i.e. NS and CCNS controllers) led to far more reliable behaviors with and without perturbation (Table 3).

Our functional analyses of the best CCNS controller give more insight into how the normalization constraint interacts with plasticity. Indeed, relaxing this constraint results in the controller inability to stop the robot when asked (Figure 12(a)). That is, synaptic

normalization seems to help stable coexistence of the two conflicting regimes underlying propulsion and movement inhibition. To explain such a multistability, we hypothesize a canalizing effect of the normalization constraint on synaptic plasticity. In particular, the induced synaptic competition might contribute to a global self-organization of individual plastic rules.

On the other hand, synaptic normalization does not seem to implicitly constrain neuronal activity as expected. If it does effectively prevent synaptic weights from saturating (Figure 12(c)), its effect on neuronal activity is not significant. In fact, when synaptic normalization is removed (Figure 12(b)), neuronal activity is left oscillating in the proper range for effective propulsion and there is not any saturated or silenced neuron.

5.3. Center-crossing constraint prevents neuronal saturation, but not robustness

In contrast to synaptic normalization, the center-crossing constraint obviously contributes to the sensitivity of the best CCNS controller. Indeed, suppressing this

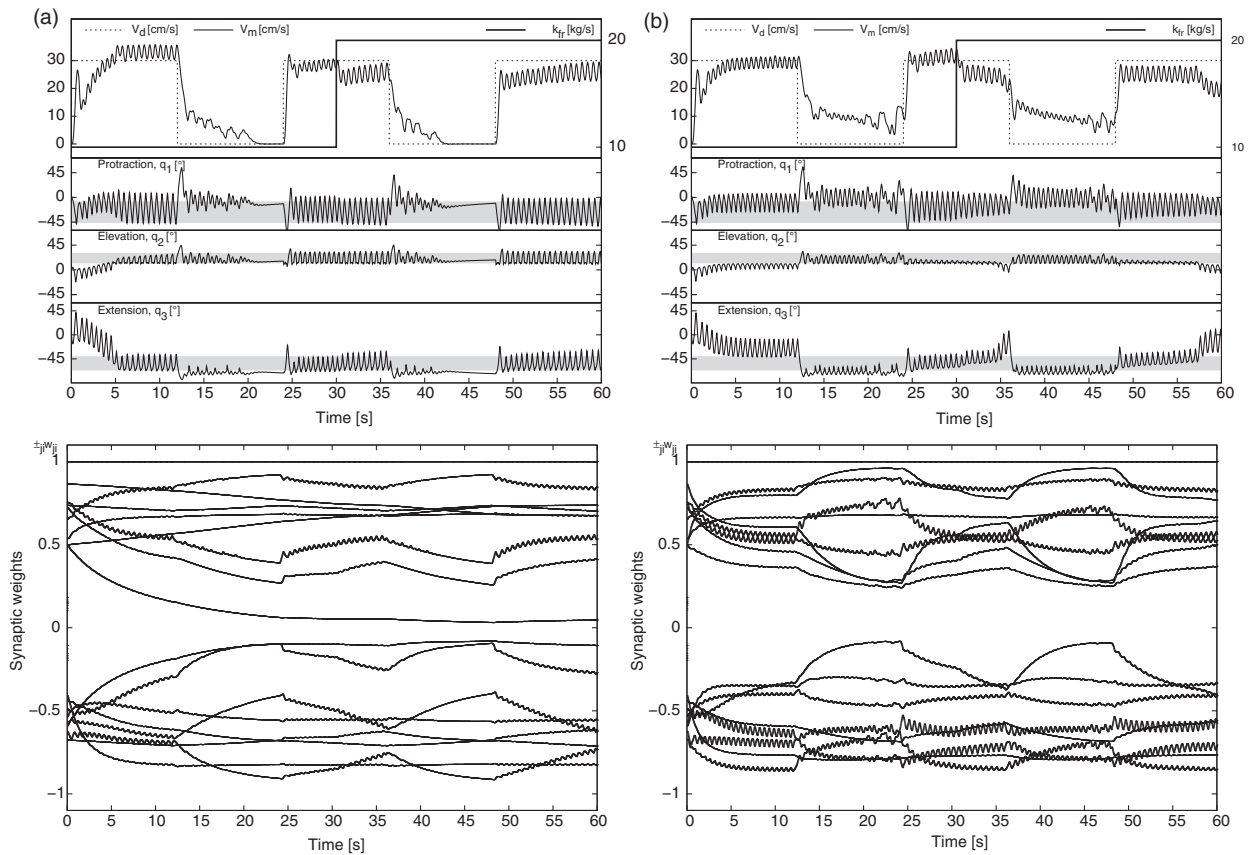


Figure 13. Effects of disturbing plasticity dynamics: (a) Tenfold multiplication of the plasticity time constants τ_{s_j} : Behavior transitions are made slower, whereas asymptotic behaviors remain unchanged. (b) All the plasticity time constants set to their mean value $\bar{\tau}_s = 1.56$ s: The controller is unable to stop the robot when needed. Propulsion behavior is also altered and seems unstable as noticed after $t = 35$ s and $t = 58$ s.

constraint leads to neuronal activity saturation, and the controller becomes fully paralyzed (Figure 12(d)).

However, the center-crossing constraint does not appear to promote the evolution of robust behaviors. Thus, as shown in our comparative analyses, performance of CC controllers drops critically in long lasting simulations when friction perturbation is triggered (Table 3). So, the enhanced evolvability of CC controllers likely comes from focusing evolutionary search on more sensitive and/or oscillating neural networks (Mathayomchan & Beer, 2002).

5.4. Robustness through homeostatic hebbian plasticity

As shown by the quick recovery from synaptic randomization perturbations (Figure 14), dynamics of the best CCNS controller is governed by very stable dynamic attractors, which are shaped by each specific plastic rule and rate tuned at every synapse (Figure 13(b)). Furthermore, in each condition of desired speed and

friction coefficient, the dynamics perturbed by random synaptic values return to the *same* previous steady state.

Therefore, the best CCNS controller clearly demonstrates the property of homeostasis as defined in the background section (Kitano, 2007). Moreover, this homeostasis underlies behavioral robustness, as it results in maintaining the evolved locomotor performance. Interestingly, homeostasis is here supported by hebbian plasticity, which is often considered destabilizing as being of positive-feedback nature. This may be due to the multistability, canalizing and sensitivity properties provided by the static constraints we proposed.

Another interesting point is that, compared to previous works, homeostasis evolved implicitly without relying on any active regulation mechanism explicitly homeostatic. Hence, it blends the distinction between homeostatic and non homeostatic mechanisms, leaving this property for rather characterizing controller dynamics.

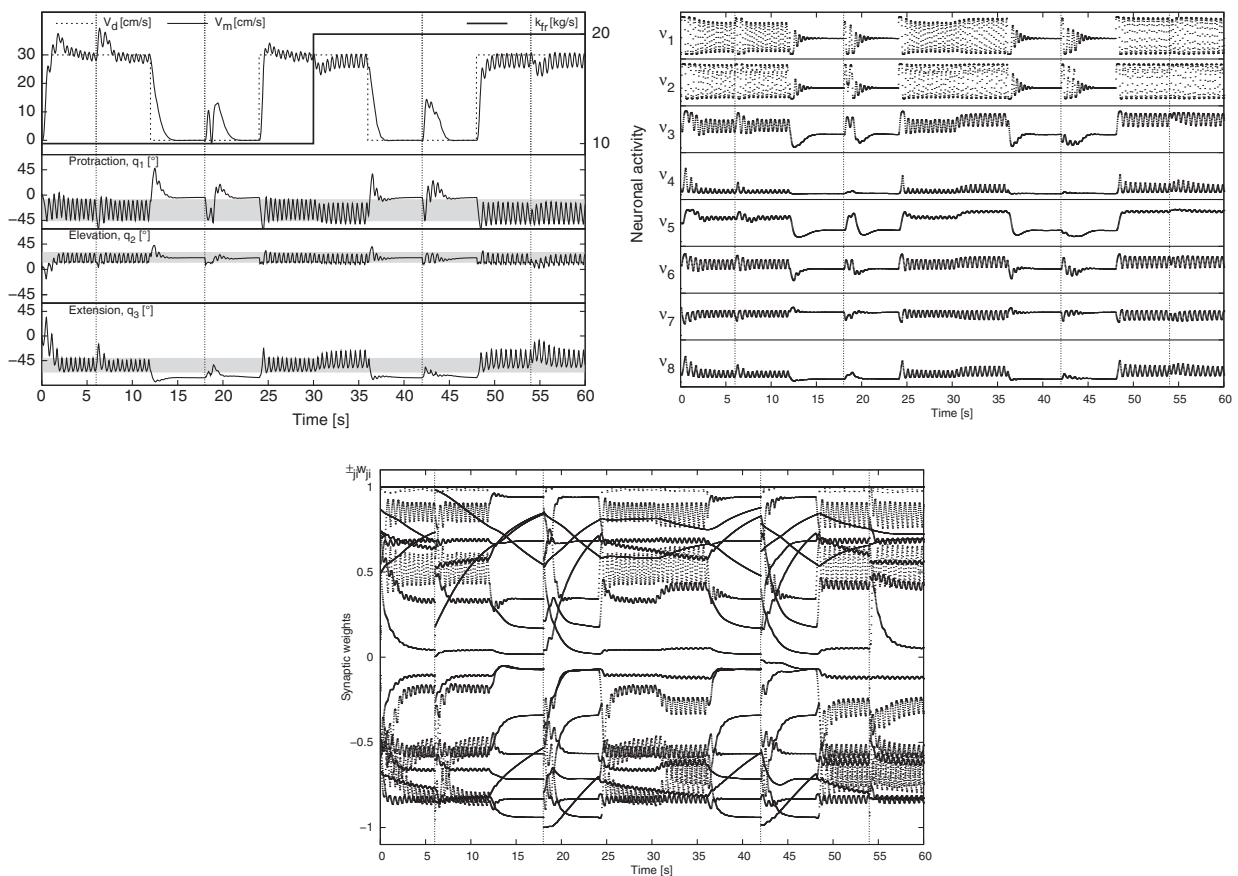


Figure 14. Robustness and homeostasis of the best CCNS controller: All synaptic weights were set to new random values, at $t = 6$ s, 18 s, 42 s and 54 s. Whatever the behavior phase in which we triggered this perturbation, the controller quickly got back to the same steady state (both activity and synapses), thus recovering the performance level.

5.5. Task-independence

So far, our findings are valid for the single-legged locomotion task only. However, we would assume that our models could be used for other problems. By being based on CTRNNs and adaptive synapses, our model should be practicable for evolving behaviors that have already been tackled in former investigations. Moreover, the proposed constraints are not related to locomotion, nor to any other task. Admittedly, more results are required to better define the area of applicability of the model.

6. Conclusion

In this paper, we have shown that plastic neurocontrollers based on adaptive synapses are more evolvable when subjected to static parameter constraints inspired from homeostatic processes of biological neurons. In particular, we could evolve both flexible and multi-stable pattern generators, leading to both adaptive and robust single-legged robot locomotion. Moreover, we

have seen that homeostasis can evolve implicitly without any active homeostatic mechanisms and be implemented through constrained hebbian plasticity.

As stable flexibility is required in many interesting behaviors, more investigations in constraining neural network models seem promising. This is especially true for plastic models which hold more potential for adaptivity but also for instability.

Further progress could be achieved through mathematical analyses of model parameter spaces (Mathayomchan & Beer, 2002; Beer, 2006; Beer & Daniels, 2010; Prinz, Bucher, & Marder, 2004; Achard & De Schutter, 2006). Unfortunately, for some complex models this might not be practical. Another approach, that we followed here, is to take inspiration from how biological neural networks tune their parameters. However, although we know developmental process and homeostatic plasticity are involved, a global understanding is not yet available (Marder, Tobin, & Grashow, 2007; Goaillard, Taylor, Schulz, & Marder, 2009). Then, how to abstract these experience-dependent processes in static parameter constraints? For a given task, which constraints would promote to evolve

homeostasis, multistability, robustness and adaptivity? A comparison between static constraints and active homeostatic mechanisms would be very helpful.

Notes

1. For a CTRNN defined with an odd activation function $\sigma(y_i^j + \theta_i)$, the center-crossing condition is verified if there is an equilibrium point for which the activations of all neurons are zero (Beer, 1995; Mathayomchan & Beer, 2002). Since $\sigma(0) = 0$, this condition occurs when $y_i^j = -\theta_i$. After substitution into equation 1 and solving for θ_i in steady state (with $I_i^j = 0$), we obtain that the center-crossing condition is given by $\theta_i^{**} = 0$.
2. Rigid body dynamics simulation has been carried out using the open source, high performance library *Open Dynamics Engine* (opende.sourceforge.net).
3. Each angle was set by applying a torque computed through a PD control loop.

Funding

This research received no specific grant from any funding agency in the public, commercial, or not-for-profit sectors.

References

Achard, P., & De Schutter, E. (2006). Complex parameter landscape for a complex neuron model. *PLoS Computational Biology*, 2(7), e94.

Baker, J. E. (1987). Reducing bias and inefficiency in the selection algorithm. In J. J. Grefenstette (Ed.), *Proceedings of the International Conference on Genetic Algorithms and their Application* (pp. 14–21). Hillsdale, NJ: Lawrence Erlbaum Associates.

Beer, R. D. (1995). On the dynamics of small continuous-time recurrent neural networks. *Adaptive Behavior*, 3(4), 469–509.

Beer, R. D. (2006). Parameter space structure of continuous-time recurrent neural networks. *Neural Computation*, 18, 3009–3051.

Beer, R. D., & Daniels, B. C. (2010). *Saturation probabilities of continuous-time sigmoidal networks*. (arXiv preprint 1010.1714).

Bienenstock, E. L., Cooper, L. N., & Munro, P. W. (1982). Theory for the development of neuron selectivity: orientation specificity and binocular interaction in visual cortex. *Journal of Neuroscience*, 2(1), 32–48.

Bullock, S. (2006). The fallacy of general purpose bio-inspired computing. In L. M. Rocha, L. S. Yaeger, M. A. Bedau, D. Floreano, R. L. Goldstone & A. Vespignani (Eds.), *Artificial Life X : Proceedings of the International Conference on the Simulation and Synthesis of Living Systems* (pp. 540–545). Cambridge, MA: MIT Press.

Di Paolo, E. A. (2000). Homeostatic adaptation to inversion of the visual field and other sensorimotor disruptions. In J.-A. Meyer, A. Berthoz, D. Floreano, H. L. Roitblat & S. W. Wilson (Eds.), *From Animals to Animats 6: Proceedings of the International Conference on the Simulation of Adaptive Behavior, SAB2000* (pp. 440–449). Cambridge, MA: MIT Press.

Di Paolo, E. A. (2002a). *Evolving robust robots using homeostatic oscillators (Cognitive Science Research Paper No. 548)*. Brighton, UK: School of Cognitive and Computing Sciences, University of Sussex.

Di Paolo, E. A. (2002b). Spike-timing dependent plasticity for evolved robots. *Adaptive Behavior*, 10(3–4), 243–263.

Di Paolo, E. A. (2003). Evolving spike-timing-dependent plasticity for single-trial learning in robots. *Philosophical Transactions of the Royal Society of London A*, 361(1811), 2299–2319.

Doncioux, S. (2003). *Évolution de contrôleurs neuronaux pour animats volants : méthodologie et applications. Doctoral dissertation in French, AnimatLab/LIP6* (pp. 149–150). Paris: Université Pierre et Marie Curie.

Floreano, D., & Mondada, F. (1996). Evolution of plastic neurocontrollers for situated agents. In P. Maes, M. J. Mataric, J.-A. Meyer, J. Pollack & S. W. Wilson (Eds.), *From Animals to Animats 4: Proceedings of the International Conference on Simulation of Adaptive Behavior, SAB1996* (pp. 402–410). Cambridge, MA: MIT Press.

Floreano, D., & Urzelai, J. (2000). Evolutionary robots: The next generation. In T. Gomi (Ed.), *Evolutionary Robotics III* (pp. 231–266). Ontario, Canada: AAI Books.

Goaillard, J. M., Taylor, A. L., Schulz, D. J., & Marder, E. (2009). Functional consequences of animal-to-animal variation in circuit parameters. *Nature Neuroscience*, 12, 1424–1430.

Goldberg, D. E. (1989). *Genetic algorithms in search, optimization, and machine learning*. Reading, MA: Addison-Wesley.

Harvey, I., Husbands, P., Cliff, D., Thompson, A., & Jakobi, N. (1996). Evolutionary robotics at Sussex. *Proceedings of the International Symposium on Robotics and Manufacturing, ISRAM1996* (Vol 6, pp. 293–298). New York: ASME Press.

Hoinville, T., & Hénaff, P. (2004a). Comparative study of two homeostatic mechanisms in evolved neural controllers for legged locomotion. *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems, IROS2004* (pp. 2624–2629). Piscataway, NJ: IEEE.

Hoinville, T., & Hénaff, P. (2004b). Evolving plastic neural controllers stabilized by homeostatic mechanisms for adaptation to a perturbation. In J. Pollack, M. A. Bedau, P. Husbands, T. Ikegami & R. A. Watson (Eds.), *Artificial Life IX: Proceedings of the International Conference on the Simulation and Synthesis of Living Systems* (pp. 81–87). Cambridge, MA: MIT Press.

Husbands, P., Smith, T. M. C., O’Shea, M., Jakobi, N., Anderson, J., & Philippides, A. O. (1998). Brains, gases and robots. In L. Niklasson, M. Boden & T. Ziemke (Eds.), *Proceedings of the International Conference on Artificial Neural Networks, ICANN98* (pp. 51–64). Berlin: Springer-Verlag.

Iizuka, H., & Di Paolo, E. A. (2007). Toward spinozist robotics: Exploring the minimal dynamics of behavioural preference. *Adaptive Behavior*, 15, 359–376.

Iizuka, H., & Di Paolo, E. A. (2008). Extended homeostatic adaptation: Improving the link between internal and behavioural stability. In M. Asada, J. C. T. Hallam, J.-

- A. Meyer & J. Tani (Eds.), *From Animals to Animats 10: Proceedings of the International Conference on the Simulation of Adaptive Behavior, SAB2008* (pp. 1–11). Berlin: Springer-Verlag.
- Jakobi, N., & Quinn, M. (1998). Some problems (and a few solutions) for open-ended evolutionary robotics. In P. Husbands & J.-A. Meyer (Eds.), *Proceedings of the European Workshop on Evolutionary Robotics, EvoRobot1998* (pp. 108–122). Berlin: Springer-Verlag.
- Kepecs, A., van Rossum, M. C., Song, S., & Tegnér, J. (2002). Spike-timing-dependent plasticity: common themes and divergent vistas. *Biological Cybernetics*, 87(5–6), 446–458.
- Kitano, H. (2007). Towards a theory of biological robustness. *Molecular Systems Biology*, 3(137), 1–7.
- Marder, E., Tobin, A. E., & Grashow, R. (2007). How tightly tuned are network parameters? Insight from computational and experimental studies in small rhythmic motor networks. *Progress in Brain Research*, 165, 193–200.
- Mathayomchan, B., & Beer, R. (2002). Center-crossing recurrent neural networks for the evolution of rhythmic behavior. *Neural Computation*, 14, 2043–2051.
- McHale, G., & Husbands, P. (2004a). GasNets and other evolvable neural networks applied to bipedal locomotion. In S. Schaal, A. J. Ijspeert, A. Billard, S. Vijayakumar, J. Hallam & J.-A. Meyer (Eds.), *From Animals to Animats 8: Proceedings of the International Conference on the Simulation of Adaptive Behavior, SAB2004* (pp. 163–172). Cambridge, MA: MIT Press.
- McHale, G., & Husbands, P. (2004b). Quadrupedal locomotion: GasNets, CTRNNs and hybrid CTRNN/PNNs compared. In J. Pollack, M. A. Bedau, P. Husbands, T. Ikegami & R. A. Watson (Eds.), *Artificial Life IX: Proceedings of the International Conference on the Simulation and Synthesis of Living Systems* (pp. 106–112). Cambridge, MA: MIT Press.
- Miller, K. D., & MacKay, D. J. C. (1994). The role of constraints in Hebbian learning. *Neural Computation*, 6, 100–126.
- Prinz, A. A., Bucher, D., & Marder, E. (2004). Similar network activity from disparate circuit parameters. *Nature Neuroscience*, 7, 1345–1352.
- Stanley, K. O., Bryant, B. D., & Miikkulainen, R. (2003). Evolving adaptive neural networks with and without adaptive synapses. *Proceedings of the Congress on Evolutionary Computation, CEC2003* (Vol. 4, pp. 2557–2564). Piscataway, NJ: IEEE.
- Tuci, E., & Quinn, M. (2003). Behavioural plasticity in autonomous agents: A comparison between two types of controller. In S. Cagnoni et al (Eds.), *Proceedings of the International Conference on Applications of Evolutionary Computing, EvoWorkshops2003* (Vol. 2661, pp. 661–672). Berlin: Springer-Verlag.
- Turrigiano, G. G., & Nelson, S. B. (2000). Hebb and homeostasis in neuronal plasticity. *Current Opinion in Neurobiology*, 10(3), 358–364.
- Urzelai, J., & Floreano, D. (2001). Evolution of adaptive synapses: Robots with fast adaptive behavior in new environments. *Evolutionary Computation*, 9(4), 495–524.
- Williams, H. (2004). Homeostatic plasticity in recurrent neural networks. In S. Schaal, A. J. Ijspeert, A. Billard, S. Vijayakumar, J. Hallam & J.-A. Meyer (Eds.), *From Animals to Animats 8: Proceedings of the International Conference on Simulation of Adaptive Behavior* (pp. 344–353). Cambridge, MA: MIT Press.
- Williams, H. (2005). Homeostatic plasticity improves continuous-time recurrent neural networks as a behavioural substrate. *Proceedings of the International Symposium on Adaptive Motion in Animals and Machines, AMAM2005*. Ilmenau, Germany: Technische Universität.
- Williams, H. (2006). Homeostatic adaptive networks. *Doctoral dissertation*. UK: University of Leeds.
- Williams, H., & Noble, J. (2007). Homeostatic plasticity improves signal propagation in continuous-time recurrent neural networks. *Biosystems*, 87(2-3), 252–259.
- Wood, R., & Di Paolo, E. A. (2007). New models for old questions: Evolutionary robotics and the “A not B” error. In F. Almeida e Costa (Ed.), *Proceeding of the European Conference on Advances in Artificial Life, ECAL2007* (pp. 1141–1150). Berlin: Springer-Verlag.

About the Authors



Thierry Hoinville received an M.Sc. in artificial intelligence from Pierre et Marie Curie University (Paris 6) in 2002 and a Ph.D. in robotics from Versailles University (UVSQ) in 2007. Since 2008, he has been a postdoctoral fellow at the Italian Institute of Technology (IIT). He has been involved in projects covering aspects of bio-inspired control for adaptive legged locomotion, reconfigurable morphology for multimodal legged locomotion, and perception in virtual reality. His research interests include evolutionary robotics, adaptive behavior, neuronal modeling, dynamic simulation, and virtual reality. He is especially interested in engineering biological principles for giving life-like motion to artificial creatures. E-mail: thierry.hoinville@iit.it



Cecilia Tapia Siles received her electromechanical engineer degree in 2003 from Universidad Privada Boliviana. In 2005, she received a scholarship from the French government to study robotics, obtaining her M.Sc. degree in robotics from Ecole Centrale de Nantes on 2006. She is currently a Ph.D. student in co-direction by Pierre et Marie Curie University (Paris 6) and Istituto Italiano di Tecnologia (IIT) Genoa University. Her research is focused on adaptive locomotion for bio-inspired underwater robots. Her research interests include biologically inspired robotics and locomotion generation and control. E-mail: cecilia.tapia@iit.it.



Patrick Hénaff received an M.Sc. degree in robotics from the University of Pierre et Marie Curie (Paris 6) in 1989, and a Ph.D. in robotics in 1994 at the Paris Robotic Laboratory, University of Pierre et Marie Curie. Since 1997, he has been Associate Professor at the Institute of Technology of Cergy Pontoise. He worked as a researcher at the LISV (Systems Engineering Laboratory of the University of Versailles) until 2009. He is now researcher at ETIS lab, University of Cergy Pontoise, CNRS UMR 8051. His topics of interest concern biologically inspired control, particularly the neural network biological models dedicated to locomotion to control dynamical behaviors and to compensate certain deficiencies of legged robots and humanoids robots. E-mail: patrick.henaff@u-cergy.fr.