Self-organization of reflexive behavior from spontaneous motor activity

Biol Cybern (2013) 107:25–37
DOI 10.1007/s00422-012-0521-7

ORIGINAL PAPER

Biological Cybernetics

Self-organization of reflexive behavior from spontaneous motor activity

Hugo Gravato Marques · Farhan Imtiaz · Fumiya Iida · Rolf Pfeifer

Received: 29 February 2012 / Accepted: 11 September 2012 / Published online: 28 September 2012

Abstract In mammals, the development of reflexes is often regarded as an innate process. However, recent findings show that fetuses are endowed with favorable conditions for ontogenetic development. In this article, we hypothesize that the circuitry of at least some mammalian reflexes can be self-organized from the sensory and motor interactions brought forth in a musculoskeletal system. We focus mainly on three reflexes: the myotatic reflex, the reciprocal inhibition reflex, and the reverse myotatic reflex. To test our hypothesis, we conducted a set of experiments on a simulated musculoskeletal system using pairs of agonist and antagonist muscles. The reflex connectivity is obtained by producing spontaneous motor activity in each muscle and by correlating the resulting sensor and motor signals. Our results show that, under biologically plausible conditions, the reflex circuitry thus obtained is consistent with that identified in relation to the analogous mammalian reflexes. In addition, they show that the reflex connectivity obtained depends on the morphology of the musculoskeletal system as well as on the environment that it is embedded in.

Keywords Self-organization · Spinal reflexes · Spinal control · Motor control · Ontogenetic development

H. G. Marques (✉) · F. Imtiaz · F. Iida
ETH Zurich, Zurich, Switzerland
e-mail: hgmarques@gmail.com

F. Imtiaz
e-mail: imtiazf@student.ethz.ch

F. Iida
e-mail: fumiya.iida@mavt.ethz.ch

R. Pfeifer
University of Zurich, Zurich, Switzerland
e-mail: pfeifer@ifi.uzh.ch

1 Introduction

The human body is a complicated and high-dimensional machine, and generating even the most primitive reflexes often requires a tight coordination between many different sensor and motor modalities (Bernstein 1967). In principle, hardwiring a given behavior implies pre-establishing every connection mediating the afferent sensor inputs and the appropriate α-motoneurons. For some reflexes this connectivity might involve an extremely large number of connections. For example, the withdrawal reflex entails a significant number of connections mediating afferent inputs from mechano (and noxious) receptors and α-motoneurons belonging to different muscles (Schouenborg and Weng 1994; Levinsson et al. 1999); in theory, it is not fully understood how this complex reflex circuitry can be developed in nature.

An alternative way to set the reflex circuitry is to branch the axon connectivity along the reflex arc to sites near the appropriate targets locations, and let experience prune the unnecessary connections. In this way, the appropriate connectivity can be statistically established in the different circuits, and ontogenetic processes ensure the removal of spurious connections. Similar processes have been observed for different areas of the central nervous system (see for example, Wang et al. 1995; Sharma et al. 2000), where local neural circuitry seems to be determined by the input information. In the spinal cord, there is also evidence that experience determines local neural circuits. For example, at birth excitatory connections mediate the connectivity between Ia muscle afferent inputs (which estimate muscle length information) and α-motoneurons of the homonymous muscle as well as of the antagonist muscles. However, during development the antagonist connections disappear while the homonymous connections are maintained (Myklebust and Gottlieb 1993).

Historically, the observation of a given behavior shortly after birth was considered as sufficient to conclude that the behavior was developed independently of experience (Robinson and Brumley 2005), i.e., it was due to some genetic (or innate) processes. This was certainly the case with many of the reflex responses humans display at birth.

The question then arises, how the presence of a given behavior at birth can be explained by ontogenetic processes. Some clues in addressing this question might lie in recent findings on prenatal development. Experiments have shown that mammalian fetuses display significant motor activity (Hadders-Algra 2007), responsiveness to sensor stimulation (Smotherman and Robinson 1992) and have the capacity for learning (Smotherman and Robinson 1998); all of which are favorable conditions for the type of ontogenetic development addressed here. In addition, it has been shown that the development of the withdrawal reflex (which is present at birth) can be explained ontogenetically using the method of motor-directed somatosensory imprinting, MDSI (Petersson et al. 2003; see also Schouenborg and Weng 1994). This method uses a form of Hebbian learning (the anti-Hebbian rule) to self-organize the tactile-to-motor circuitry entailed by the withdrawal reflex.

In this article, we propose a novel framework that characterizes a few important aspects of ontogenetic development in mammals. Our framework extends the MDSI method to incorporate two additional sensor modalities—muscle length and muscle force—to obtain three additional reflexes: the myotatic reflex, the reciprocal inhibition reflex, and the reverse myotatic reflex. These reflexes consist of fast (negative) feedback responses involved in local agonist and antagonist muscle interactions and they were chosen because they are generally placed at the lowest level in mammalian motor control (Latash 2008). Moreover, they are relatively simple (they involve at most one interneuron in their reflex arc) and fairly well understood, which allows for validation.

Our hypothesis is that these three reflexes can emerge from the interactions between a mechanism capable of producing spontaneous motor activity on a musculoskeletal system and a correlation-based method applied on the resulting sensor and motor activity. In our framework, spontaneous motor activity consists of spontaneous contractions of individual muscles, or single muscle twitches (SMTs) which typically occur during sleep (Blumberg and Lucas 1996). In addition, we hypothesize that two factors strongly contribute to the normal development of the reflexes: a perturbation-free environment and the morphology of biological muscles.

This is our first systematic investigation on the self-organization of reflexes. It is not yet intended to involve accurate modelling of all human neurophysiological functions, but rather to extract more general biological principles which can also be used in the domain of artificial systems.

The remainder of this article is organized as follows. Section 2 describes our framework. Section 3 provides the implementation details of each mechanism in the framework. Section 4 describes the methods and results of five experiments: the first experiment investigates the self-organization of the three reflexes using a biological muscle model, the second investigates the role of sensor delays in the reflex connectivity, the third analyses the reflex connectivity using synergistic muscles, and the fourth and fifth experiments investigate the role of the environment and the muscle morphology, respectively, on the development of the reflex circuitry. Section 5 discusses the experimental results, and it is followed by Sect. 6 which provides the conclusions and the outlook of our research.

2 The reflex-learning framework

The proposed framework consists of five interacting models: a musculoskeletal model (and its environment), a peripheral model, a model of SMTs, a learning model based on the correlations between sensor and motor activity, and a model of the reflex sensorimotor mapping (see Fig. 1). It works as follows. First, SMTs activate the different motors of the system. Second, the activated motors produce forces which are propagated through the musculoskeletal system (and the environment that it is embedded in). Third, the changes produced in the musculoskeletal system are captured by the different body sensors, which (fourth) convert them into sensor activity. And fifth, the correlation between the sensor and motor activity establishes the reflex connectivity between each receptor and each effector. These models interact in two stages: in the first stage, the models interact to learn the reflex circuitry; in the second stage, the models interact to produce the reflex activity. In this section, we will introduce each of the five models of the framework, and in the next section, we will describe their implementation.

2.1 Musculoskeletal and environment model

The musculoskeletal and environment model captures the mechanical interactions among muscles, skeleton, and environment. When a muscle contracts,[1] it produces motion on the skeleton which induces sensory information in the system. These motions depend not only on the skeletal properties (e.g., length and mass of the bones) but also on the characteristics of the environment where it is embedded in (e.g., gravity or viscosity).

[1] Often, a muscle contraction is associated to the shortening of a muscle; here, we will use the term *muscle contraction* to mean the generation of tension in the muscle, irrespective of whether this tension causes the muscle to shorten or not.

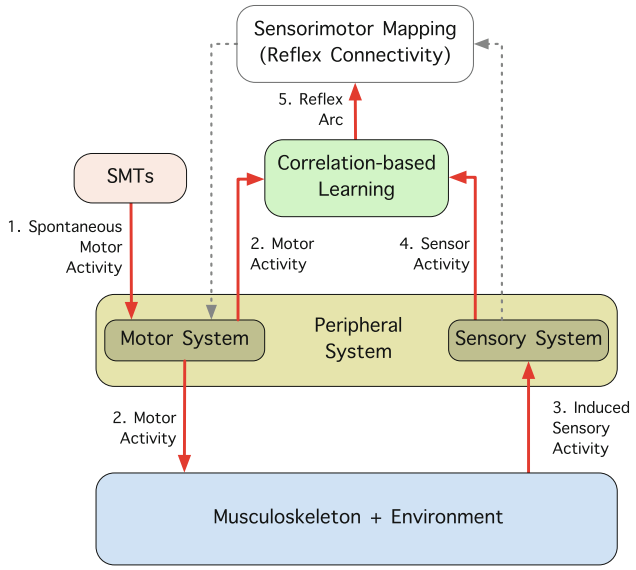

Fig. 1 The proposed framework consists of five interacting models: a musculoskeletal model (and its environment), a peripheral model, a model of SMTs, a learning model based on the correlations between sensor and motor activity, and a model of the reflex sensorimotor mapping (see text). *Numbers* The order of events in the framework

The muscles in biology have two fundamental properties. First, they have asymmetrical conditioning, i.e., they only produce active forces when contracting but not when extending. And second, they impose a very small resistance to movement when relaxed. All these properties are a part of the standard Hill-type muscle model which is often used to simulate the behavior of biological muscles (Hill 1938).

2.2 Peripheral system model

The peripheral system allows the communication between the neural system and the mechanical system (e.g., the musculoskeletal system), via sensor inputs and motor outputs. In humans, afferent fibers collect the sensory signals from the different sensory modalities, and efferent fibers carry the motor signals to the motoneurons. In the context of our framework, the relevant afferent fibers are the Ia and Ib fibers; group II afferent fibers are out of the scope of this article. The former estimate changes in the muscle length as well as positional muscle length (Hulliger 1984); the latter respond to small variations in muscle force, and fire mostly when the muscle is being contracted (Jami 1992). The relevant efferent fibers are the α-motoneurons which are responsible for initiating muscle contractions.

2.3 SMTs model

SMTs are one type of spontaneous motor activity observed in mammals (Blumberg and Lucas 1994; Blumberg and Lucas 1996); they consist of single muscle contractions which occur in the absence of sensory stimulation. In mammals, this type of activity has been observed prior to birth as well as after birth, during sleep (Robinson et al. 2000). Before birth, spontaneous motor activity has been argued to be a major driving force in the development and maturation of the nervous system (Prechtl 1990; Smotherman and Robinson 1996; Hadders-Algra 2004, 2007; Robinson and Brumley 2005). In fetuses, development is intrinsically driven by self-organization processes which structure the sensorimotor information flow in the system according to the systematic exploration of contingent sensor and motor activity (Blumberg and Lucas 1994).

Spontaneous motor activity has been argued to be an essential mechanism in reducing the dimensionality of the motor space (Sporns and Edelman 1993; see also Bernstein 1967) through the formation of orderly connectivity (Blumberg and Lucas 1996), the elimination of spurious synapses, and the pruning of muscle fibers (Blumberg and Lucas 1994).

2.4 Correlation-based learning model

Correlation-based learning methods are unsupervised-learning mechanisms that rely only on the statistical properties of their input signals. This is the feature that separates unsupervised from supervised learning methods. In previous simulation studies, it has been shown that meaningful patterns of sensorimotor activity emerge from self-generated movements induced on a physical body (e.g., Lungarella and Sporns 2005; Kuniyoshi and Sangawa 2006). More relevantly, it has also been shown that local feedback circuitry can be obtained through the self-organization of sensorimotor activity (Todorov and Ghahrammani 2003; Petersson et al. 2003; Grillner 2004; Berniker et al. 2009; Mori and Kuniyoshi 2010).

The most famous self-organization principle in neural systems is that of Hebbian learning, according to which neurons that "fire together should wire together" (Hebb 1949). The self-organization principle used here is based on the method of MDSI, which consists of the additive inverse of the temporal correlation between the sensor and motor activity. This mechanism is often referred to as anti-Hebbian (Földiák 1990) or a reversed type of Hebbian learning "with the postsynaptic activity in reflex interneurons (and motoneurons) preceding the afferent input" (Petersson et al. 2003; see also Schouenborg and Weng 1994).

2.5 Sensorimotor mapping model

When the learning phase is completed, the outcome of the correlation-based learning model is turned into a sensorimotor mapping model. More specifically, as soon as the system completes the muscle twitches, the learned connectivity is

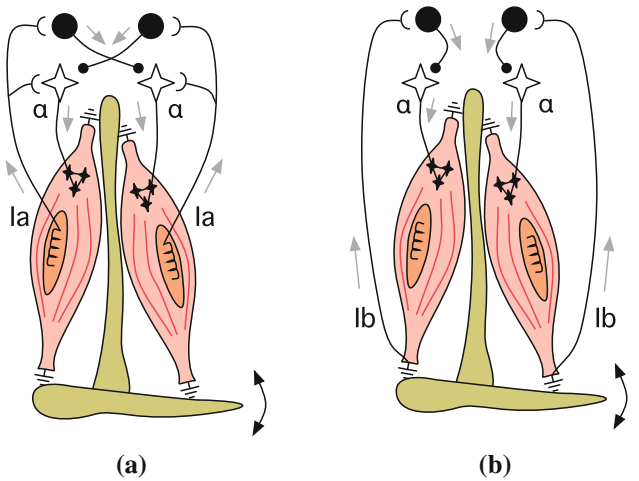

Fig. 2 The three reflexes investigated. *Stars* α-Motoneurons, *large solid circles* inhibitory interneurons, *semi-circled arcs* excitatory connections, *small solid circles* inhibitory connections. **a** The myotatic reflex is carried out through an excitatory connection between the Ia fibers and the α-motoneurons of the homonymous muscle. The reciprocal inhibition reflex is carried out through an inhibitory connection between the Ia fibers with the α-motoneurons of the antagonist muscle; this connection is mediated by an inhibitory interneuron. **b** The reversed myotatic reflex is carried out through an inhibitory connection between the Ib fibers and the α-motoneurons of the homonymous muscle; this connection is also mediated by an inhibitory interneuron

used to drive the whole system in such a way that sensory stimulation induces muscle activity based on the reflex circuitry. A description of the target reflexes is given below.

The myotatic reflex counteracts an undesired stretch imposed by an external load (Bear et al. 2007, pp. 439–440). When a load causes a muscle to stretch, it activates its Ia fibers which in turn excite the α-motoneurons of the homonymous muscle (Kudo and Yamada 1985; Chen et al. 2003; Pierrot-Deseilligny and Burke 2005, pp. 63–66). This results in a muscle contraction which counteracts the force exerted by the external load (see Fig. 2a).

The reciprocal inhibition reflex is involved in agonist–antagonist interactions; the stretch and reflexive contraction of one muscle inhibits the antagonist muscle and prevents it from counteracting the movement initiated by the agonist muscle (Pierrot-Deseilligny and Burke 2005, pp. 197–200). The combination of the myotatic reflex and the reciprocal inhibition reflex is often called the stretch reflex.

The reverse myotatic reflex is supposed to prevent muscles from producing excessive forces (Bear et al. 2007, p. 445), although its exact function is still disputed (Pierrot-Deseilligny and Burke 2005, pp. 256–257). Strong muscle contractions cause large force increases in the tendons which activate the Ib fibers. This activity excites inhibitory interneurons, which in turn inhibit the α-motoneurons of the homonymous muscle (see Fig. 2b), resulting in a decrease of the overall muscle activity.

3 Implementation of models

3.1 Implementation of musculoskeleton and environment

To simulate the interactions between two antagonist muscles, we used a virtual 3D model of two human legs. The simulation of the leg dynamics has been carried out using CALIPER (Wittmeier et al. 2011b). The framework is based on Bullet Physics[2] and Coin3D graphics,[3] and it allows to simulate in real time the interactions between a large number of rigid bodies, different types of joints and different muscle models.

The skeletal model comprises two hip joints, which are modeled as hinge joints. Each leg is a mechanically independent system, i.e., no energy is transferred from one leg to the other, to allow for an accurate validation of our learning framework. For the actuation, we use an agonist and antagonist pair of muscles in each leg: the Gluteus Maximus and the Illiacus, which are mainly responsible for hip extension and flexion, respectively (see Fig. 3). Two additional muscles, Ext.1 and Ext.2, have been implemented in each leg; these muscles are used only to test the reflex circuitry (see Sect. 4.2) and to test the impact of external perturbations on the reflex connectivity (see Sect. 4.5). When referring to one of these muscles, we will explicitly use the term *external muscle*; unless stated otherwise, the term *muscle* will be only used to refer to the leg muscles (i.e., the Illiacus or the Gluteus Maximus).

Each muscle is simulated as a straight line between two rigid bodies (see Fig. 3). The main muscle model used in our investigation is based on a two-element nonlinear Hill model (Hill 1938; Zajac 1989). The two elements are an active contractile element in parallel with a passive elastic element (see Fig. 4a). The contractile element models the active force generated by the muscle fibers. This element includes a damping mechanism that simulates the force–velocity relation of biological muscles. The passive elastic element models the muscle fiber's resistance to deflection and prevents the muscle from getting slack. The force produced by muscle i at its attachment points is given by:

$$F_{\mathrm{M}_i} = F_{\mathrm{CH}_i} + F_{\mathrm{SH}_i}, \tag{1}$$

where F_{CH_i} is the force produced by the Hill contractile element of muscle i and F_{SH_i} is the force produced by the passive spring element of muscle i. These forces are given by:

$$F_{\mathrm{CH}_i} = \frac{M_i}{1 + C_{\mathrm{H}} \dot{l}_i^2} \tag{2}$$

[2] Bullet Physics Library. Sony Computer Entertainment. Available online: http://www.bulletphysics.com.

[3] Coin3D. Kongsberg Oil and Gas Technologies. Available online: http://www.coin3d.org.

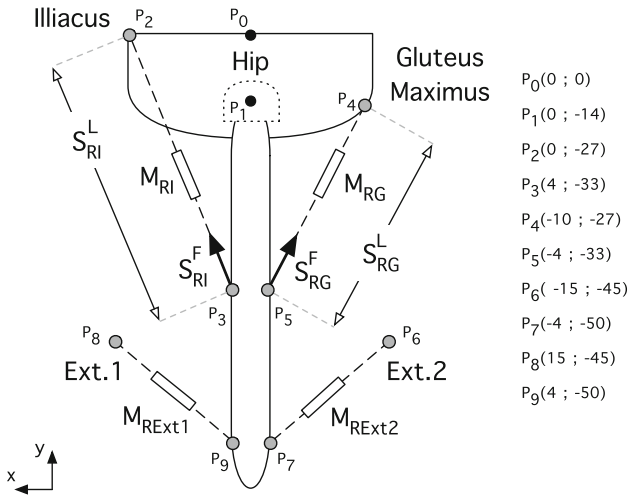

Fig. 3 Diagram of the implemented leg model with the projection of the right leg onto the sagittal plane. This same model is used for the right and left legs. Each muscle consists of a straight line (represented with *dashes*) connecting two attachment points (*filled circles*). Each leg has four muscles, M: the Illiacus, the Gluteus Maximus, and the external muscles, Ext.1 and Ext.2. Each muscle has two sensors: one that estimates the force, S^F, produced at the attachment points, and one that estimates the length, S^L, of the muscle. The subscripts RI and RG stand for the right Illiacus and right Gluteus Maximus, respectively. The coordinates on the *right* show the position of all relevant elements in the model (in centimeters)

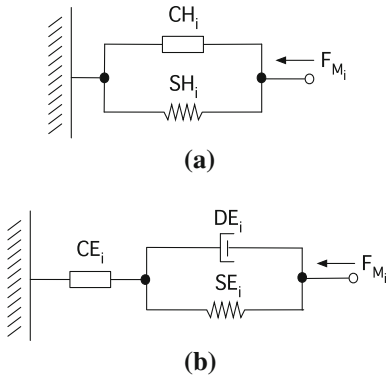

Fig. 4 The schematics of the two muscle models used in this article. **a** The Hill muscle model and **b** the artificial muscle model. Both models include a contractile element (represented by a *square*) which produces a force that brings the extremities of the muscle together. Spring and damper elements are represented by their standard symbols

$$F_{\mathrm{SH}_i} = K_{\mathrm{H}}(l_i - l_{\mathrm{r}}), \tag{3}$$

where C_{H} and K_{H} are constant factors, M_i is the motor activation, l_i is current length of muscle i, $\dot{l}_i$ is the rate of the muscle length change, and l_{r} is resting length of the muscle.

In biology, the force generated by the passive spring element of the muscle, F_{SH_i}, is significantly smaller than the force generated by the contractile element, F_{CH_i}. In our model, we set $K_{\mathrm{H}} = 1$ and $C_{\mathrm{H}} = 1e6$, which fit with the general muscle properties stated above. We have varied these parameters within one order of magnitude and no significant effect was observed on our results.

To compare the impact of the muscle model on the reflex connectivity (see Sect. 4.6), we implemented a second muscle model inspired in the artificial muscle model used in the anthropomimetic robot, ECCEROBOT (Holland and Knight 2006; Wittmeier et al. 2011a). The artificial muscle model simulates the dynamics of a DC motor with a gearbox in series with a cable and an elastic shock cord (see Fig. 4b). When the motor is actuated it reels the cable and expands the shock cord; this produces a force that brings the attachment points closer and simulates a muscle contraction.

The main reason to include an artificial muscle model in our investigation originates from the fact that DC motors behave differently from biological muscles; namely, they are not back-drivable which prevents them from passively extending when relaxed.

The force, F_{M_i}, produced by the artificial muscle is given by:

$$F_{\mathrm{M}_i} = F_{\mathrm{CE}_i} = F_{\mathrm{SE}_i} + F_{\mathrm{DE}_i} = K_{\mathrm{E}} \cdot (l_i - l_{\mathrm{r}}) \tag{4}$$

$$+ D_{\mathrm{E}} \cdot \dot{l}_i = \frac{\tau_i}{r}, \tag{5}$$

where F_{SE_i} is the force produced by the artificial contractile element, F_{S_i} is the spring force in the kiteline of muscle i, F_{DE_i} is a damping force of muscle i, K_{E} is the spring constant, D_{E} is the damping constant, r is the radius of the motor shaft, and τ_i is the motor torque at the gearbox output shaft in the artificial muscle i. This toque is calculated as a function of the motor activation, M_i:

$$\tau_i = f(M_i), \tag{6}$$

where f simulates the interactions of an electric motor and a gearbox. As the muscle can only produce forces when contracting but not when extending, the additional condition $F_{\mathrm{M}_i} \geq 0$ is added for all the artificial muscles. The details of this model as well as its parameters can be found in Wittmeier et al. (2011a).

3.2 Implementation of peripheral system

Both muscle models include two types of sensors: one that measures the length of the muscle, S^L (i.e., the distance between the two attachment points), and one that measures the force at the attachment points, S^F; when referring to sensors collectively, we will simply use the symbol S. In our simulation, the sensor activity consists of the derivative values of these sensor inputs, and the motor activity consists of the motor activations, M_i, of all the muscles (see Sect. 2.2).

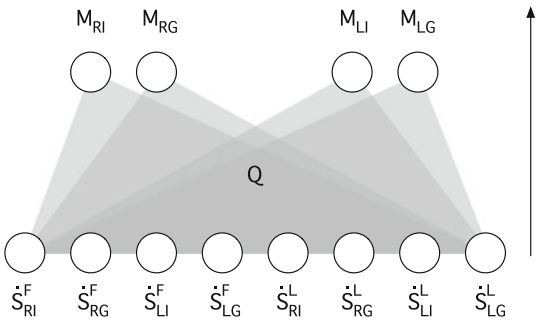

Fig. 5 The reflex network before the self-organizing process; initially the network is fully connected. *Circles* Sensor or motor elements, *shaded areas* all possible connections between sensor and motor elements, *arrow* the direction of the connectivity, M stands for motor, $\dot{S}$ stands for sensor derivative. The superscripts, F and L, stand for force and length sensors, respectively. The subscripts, I and G, stand for Illiacus and Gluteus Maximus, respectively, and the subscripts L and R stand for left and right legs, respectively. The weight matrix, Q, represents the network connectivity matrix (see text)

3.3 Implementation of correlation-based learning

The self-organization process uses a variant of differential Hebbian learning (Kosko 1986), which utilizes the anti-Hebbian rule. The process basically computes a negative cross-correlation between motor and sensor signals. All possible combinations of sensor and motor connections are considered (see Fig. 5). The connectivity $Q_{i,j}$ between motor element i and sensor element j is given by:

$$Q_{i,j} = -\eta_{ij} \sum_{t=1}^{T} M_{i,t} \cdot \dot{S}_{j,t-D} \tag{7}$$

$$\eta_{ij} = \frac{1}{\max(\dot{S}_j) \sum_{t=1}^{T} M_{i,t}}, \tag{8}$$

where $\dot{S}_j$ is the derivative of sensor j, M_i is the signal sent to motor i, η_{ij} is a normalization factor, and D is the time delay in the sensor response. Excitatory connections are characterized by positive values and inhibitory connections by negative values. The strength of each connection is given by its magnitude. Initially, all the elements in Q are set to zero.

3.4 Implementation of spontaneous motor activity

In the simulation mode, SMTs are generated as follows. First, at a given time the system randomly selects a muscle from a uniform distribution. Second, the system activates the selected muscle with a constant value ($M = 2$ mu)[4] for a fixed period of time ($t = 1.0$ s). Third, the system waits for a another period of time before triggering the next twitch. During this period all the muscles are relaxed ($M = 0$ mu). A waiting period of $t = 20$ s was found large enough for the system to recover from the previous twitch.

[4] where mu stands for muscle units. Note that mu is used here as the measurement unit used to measure the motor activity, and it is not to be confused with the motor units existing in a muscle.

3.5 Implementation of sensorimotor mapping

Intuitively, the connectivity in Q describes motor-to-sensor connections, as the directed flow of information is from motors to sensors. However, Q can also describe directed sensor-to-motor connectivity (see Petersson et al. 2003).

This operation can be technically achieved by computing the motor activities from the sensory signals based on the following equation:

$$M_i = \sum_{j=1}^{m} Q_{i,j} \frac{\dot{S}_j}{\max(\dot{S}_j)}, \tag{9}$$

where m is the number of motors in the system.

To allow for a stable simulation, we compute the motor activity based on an average of five samples of the sensor derivative signals. The collection of the five samples starts as soon as one sensor derivative is identified that goes above a given threshold. The motor signal is kept constant for the next five samples, time at which the previous five sensor readings are averaged again and a new motor signal is computed.

To produce the external stimuli, we contract the external muscles, Ext.1 and Ext.2, in each leg. Each of these muscles produces a given displacement in one of the limbs which is detected by the different sensors and converted into reflex motor activity using the matrix Q.

4 Experiments

This section describes five experiments. The first experiment verifies our main hypothesis, i.e., that our framework is capable of self-organizing the three reflexes described above. The second experiment demonstrates that sensor delays need to be significantly large to influence qualitatively the reflex connectivity. The third experiment uses a leg model with eight muscles to show the scalability of our framework as well as its potential for dimensionality reduction. The fourth experiment investigates the role of external perturbations during learning. And the fifth experiment investigates the impact of using the artificial muscle model on the reflex circuitry obtained. Before describing the experiments, we provide the initial conditions for all the experiments.

4.1 Initial conditions

All experiments follow the same methodology. They start with all the muscles relaxed ($M = 0$ mu). In this condition, the legs fall straight down due to the effects of gravity. All the connections in Q are initially set to zero. A series of muscle twitches are then carried out while sensor and motor data are being collected.

4.2 Experiment 1: self-organization of basic reflexes

4.2.1 Methods

The main goal of Experiment 1 is to verify that in normal conditions the reflex connectivity can be learned using SMTs and the self-organization process mentioned above. The Hill-type muscle model is used. No time delay is considered between sensor and motor elements ($D = 0$). This experiment is divided into two stages: the learning and the testing stages. In the learning stage, the system learns the connectivity matrix Q using the SMTs. In the testing stage, we test the reflex behavior by contracting each one of the external muscles and observe how the system reacts to the induced sensory stimulation.

4.2.2 Results

The raw sensor and motor activity generated during the learning stage is shown in Figs. 6, 7, and 8. Figure 6 shows the raw motor activity produced by the SMTs. As can be seen, at any given time only one motor at most is active. Figure 7 shows the raw length data collected for the two muscles of the right leg during a SMT carried out by the Illiacus (a) and by the Gluteus Maximus (b). As can be observed, the length sensors change their values in response to contractions of both the homonymous as well as the antagonist muscles.

Figure 8 shows the raw force data collected for the two muscles of the right leg during a SMT carried out by the Illiacus and by the Gluteus Maximus. In contrast to the length sensors the activity of the force sensors is only observed during contractions of the homonymous muscle but not during contractions of the antagonist muscle. This is because when a muscle is relaxed the only force in the muscle is due to the passive spring element which has a negligible magnitude when compared with the active force that can be produced by the contractile element of the muscle.

The data collected during the learning stage suggests a causal relationship between the force sensors and their homonymous motors, and a causal relationship between the length sensors and both their homonymous and antagonist muscles. Figure 9 shows that this is in fact the connectivity obtained. As expected, no connectivity has been created between sen-

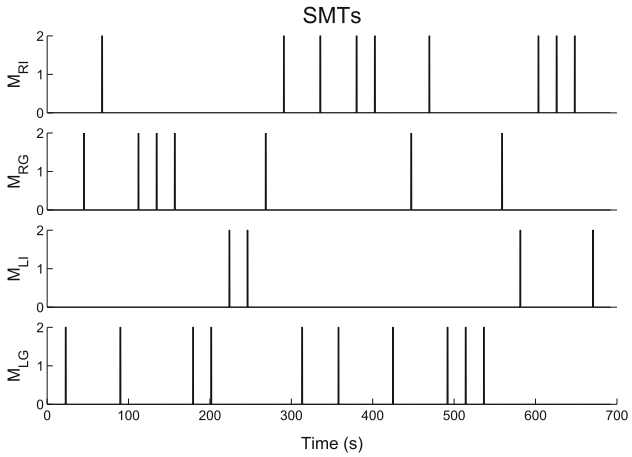

Fig. 6 Motor activity produced by SMTs; the motor activity is given in mu (motor units)

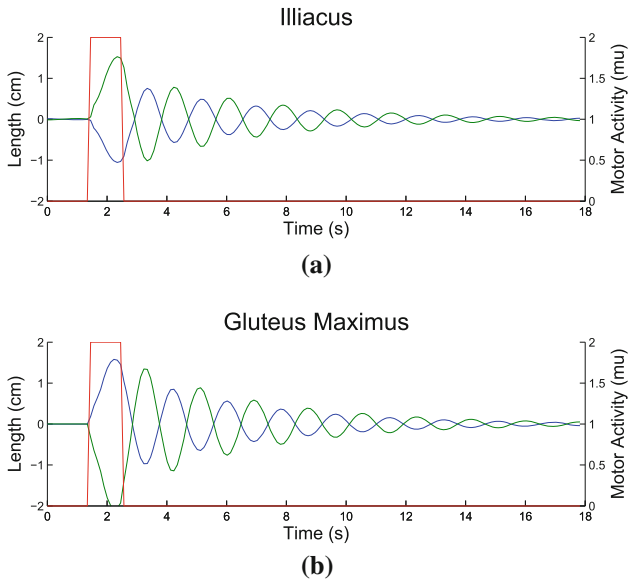

Fig. 7 The activity of the length sensors relative to the Illiacus (*blue*) and Gluteus Maximus (*green*) of the right leg in response to a SMT (*red*) carried out by **a** the Illiacus and **b** the Gluteus Maximus. For clarity, we have subtracted the resting lengths to the total length of each muscle. Similar data have been observed for the left leg. (Color figure online)

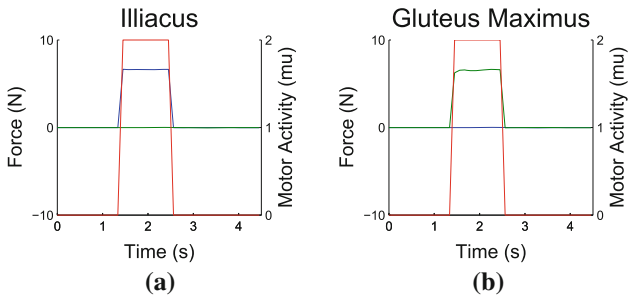

Fig. 8 The activity of the force sensors relative to the Illiacus (*blue*) and Gluteus Maximus (*green*) of the right leg in response to a SMT (*red*) carried out by **a** the Illiacus and **b** the Gluteus Maximus. Similar data have been observed for the left leg. (Color figure online)

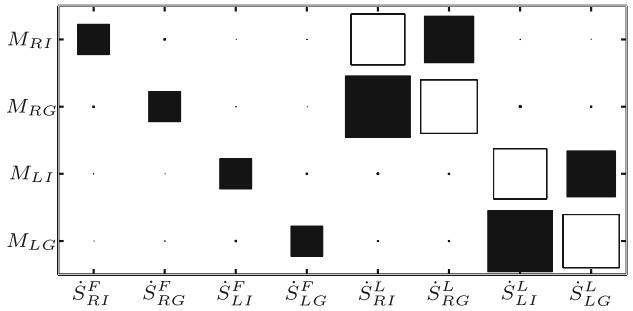

Fig. 9 The adjacency matrix, Q, obtained in Experiment 1. *Black blocks* Inhibitory connections, *white blocks* excitatory connections

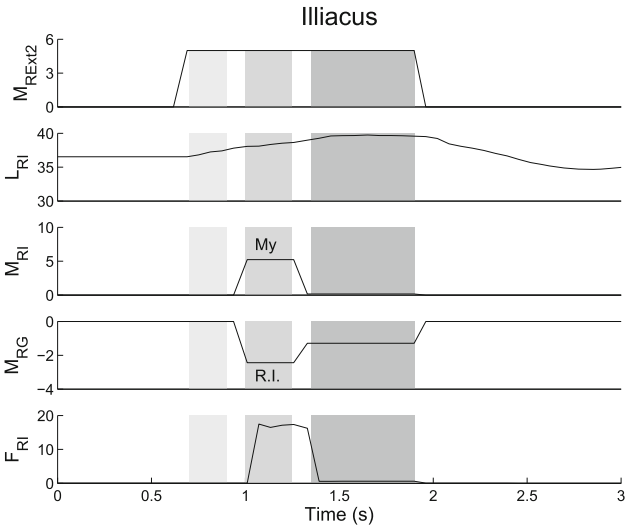

Fig. 10 Muscle activity generated in reaction to a stretch in the Illiacus caused by an external load imposed by the right Ext.2. From *top to bottom*, (1) the motor activity of the right Ext.2, (2) the length of the right Illiacus, (3) the motor activity of the right Illiacus, (4) the motor activity of the right Gluteus Maximus, and (5) the force of the right Illiacus. My refers to motor activity caused by the myotatic reflex, and R.I. refers to motor activity caused by the reciprocal inhibition reflex. *Light shaded area* The time between when the load has been initiated and the activation of the reflexes, *medium shaded area* the activation of the myotatic and reciprocal inhibition reflexes, and *dark shaded area* the motor activity after triggering these reflexes

sors from one leg and muscles from the other (as the muscles of one leg do not induce sensory stimulation on the other).

The connectivity matrix is shown in Fig. 9. The connectivity obtained is in qualitative terms similar to that observed in relation to the human reflexes (see Fig. 2). First, we obtain excitatory connections between the length sensors and their homonymous muscles as in the myotatic reflex (see for example the connection between M_{RI} and $\dot{S}^L_{\mathrm{RI}}$). Second, we obtain inhibitory connections between the length sensors and their antagonist muscles as in the reciprocal inhibition reflex (see for example the connection between M_{RI} and $\dot{S}^L_{\mathrm{RG}}$). And third, we obtain inhibitory connections between the force sensors and their homonymous muscles as in the reverse myotatic reflex (see for example the connection between M_{RI} and $\dot{S}^F_{\mathrm{RI}}$). The lower weight obtained with the force sensors is justified by the fact that significant changes in force are only observed when the twitch starts, which decreases the value of the correlation for the entire duration of the muscle twitch.

A number of marginal connections can also be seen in Fig. 9; these connections are shown as small dots (e.g., M_{LG} and $\dot{S}^L_{\mathrm{RI}}$). The weights of these connections are three orders of magnitude lower than those relative to the three reflexes, and can therefore be safely neglected.

Though the main goal of this experiment is to analyze the reflex connectivity obtained using our framework, we also tested the corresponding muscle activity in response to a muscle stretch caused by the external muscles. The testing of the reflex behavior (during the testing stage) is shown in Fig. 10. After the load is applied (light shaded area), the resulting muscle stretch causes the Illiacus to contract (medium shaded area), and at the same time inhibits the antagonist muscle—the Gluteus Maximus—from contracting (medium shaded area). This is consistent with the myotatic and reciprocal inhibition reflexes, respectively. During this first activation period, the rate of length change of the Illiacus decreases (medium shaded area) and the force in the muscle increases (medium shaded area). This causes a drop in the subsequent motor activity of the Illiacus, which is accompanied by a decrease in the inhibition of the Gluteus Maximus (dark shaded area). Qualitatively, this behavior is consistent with that observed relative to the analogous human reflexes.

To investigate the role of the reverse myotatic in the decrease of muscle activity, we repeated the experiment with and without the force component, and compared the muscle activations in both conditions. We observed that the decrease in muscle activity is mainly due to the decrease of the muscle velocity (detected in the length sensor) rather than due to the reverse myotatic reflex. The drop due to the presence of the force component is in fact very small—0.25 mu (see Sect. 5.3 for discussion).

4.3 Experiment 2: reflex connectivity with sensor delays

4.3.1 Methods

The absence of sensor signal delays in Experiment 1 is unrealistic; this is coped with in the experiment described in this section. The goal of Experiment 2 is to evaluate the extent to which sensor delays affect the reflex connectivity obtained. In this experiment, we basically evaluate the data collected for Experiment 1 but we introduce time delays in the sensor input.

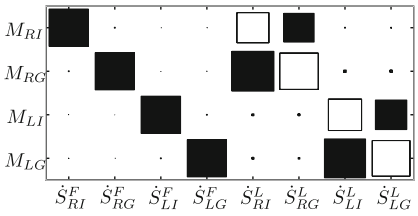

(a) $D = 0.8s$

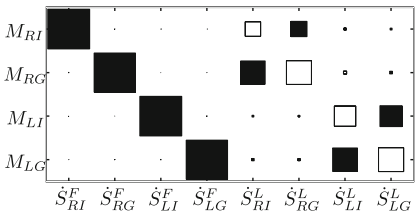

(b) $D = 0.9s$

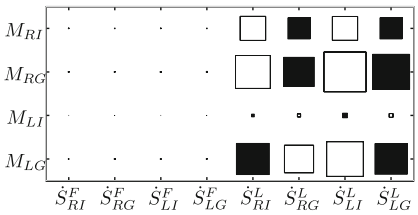

(c) $D = 1s$

Fig. 11 Adjacency matrices, Q, obtained for time delays: **a** $D = 0.8$ s, **b** $D = 0.9$ s, and **c** $D = 1$ s

4.3.2 Results

Figure 11 shows the connectivity obtained for three different delays. As can be observed the appropriate reflex connectivity is only broken at $D = 1$ s. Our framework is made to capture overlaps between sensor and motor activity. Given that at $D = 1$ s, the delay reaches the duration of the muscle twitch (1 s), we observe no overlap between sensor and motor activity, and therefore no meaningful correlations are extracted.

In addition, we observe that the relative magnitude of the connections changes with the sensor delay (compare Figs. 9 and 11a, b). We observe that the magnitude of the connections with length sensors decreases relative to the connections with force sensors. This is because the only meaningful activity in the force sensors occurs only at the very beginning of a twitch (when a force increase is observed). This means that for lags smaller than the duration of the twitch the correlation values remain basically unchanged. This contrasts with the length sensors, the magnitude of which increases steadily throughout the entire duration of the twitch. This makes the magnitude of the correlation with these sensors more sensitive to delays.

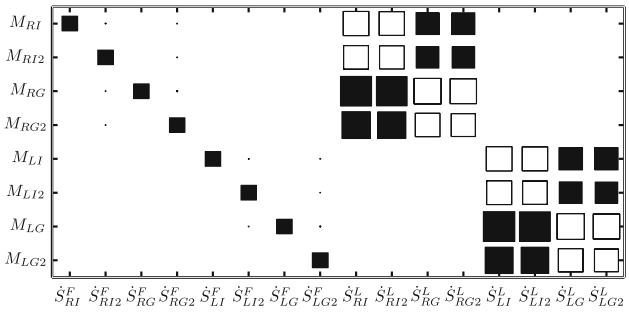

Fig. 12 Adjacency matrix, Q, obtained using eight muscles

4.4 Experiment 3: reflex connectivity with eight muscles

4.4.1 Methods

The goal of Experiment 3 is to show that our method scales to more complex systems. In particular, this experiment considers a system which, in addition to homonymous and antagonist interactions, includes synergistic interactions (i.e., muscles which produce similar motions at the joint). The success of this experiment serves to strengthen our claim that our framework can be used for the purposes of dimensionality reduction. For this purpose, we conducted experiments similar to Experiment 1, but here we implement an additional synergistic muscle paired to each of the muscles in Fig. 3. Each of the synergistic muscles is labeled by adding the number "2" to the original muscle; for example, the synergistic muscle of RI is labeled RI2. The attachment points of each of the synergistic muscles are located 5 cm below the attachment points of the original muscle. Because, we now have twice as many muscles, we carry out 60 SMTs during training.

4.4.2 Results

Figure 12 shows the connectivity obtained in Experiment 3. As in Experiment 1, connections with force sensors are only identified within the same muscle. This is expected as the muscle only presents a meaningful force change when it is (itself) active. More relevant are the connections identified with the length sensors. The connections between motor elements and length sensors belonging to synergist muscles (e.g., M_{RI} and $\dot{S}^L_{\mathrm{RI2}}$) are consistently excitatory, whereas connections between motor elements and length sensors belonging to antagonist muscles are consistently inhibitory (e.g., between M_{RG} and $\dot{S}^L_{\mathrm{RI}}$, or between M_{RG} and $\dot{S}^L_{\mathrm{RI2}}$). This means that the connectivity identified is capable of responding to a stretch in a given muscle by recruiting automatically all its synergistic muscles and inhibiting all its antagonists.

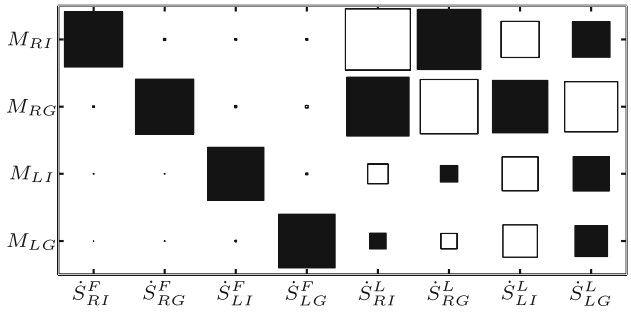

Fig. 13 Adjacency matrix, Q obtained for a perturbation probability of 0.2

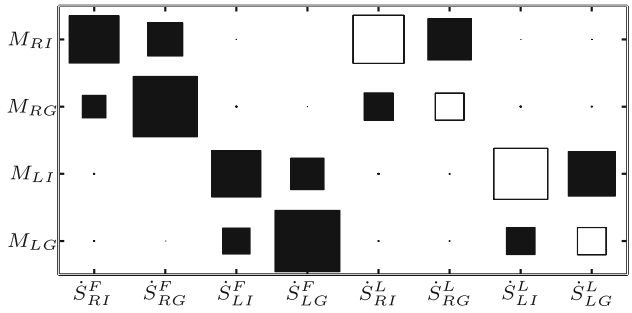

Fig. 14 Adjacency matrix, Q, obtained using the artificial muscle model

4.5 Experiment 4: the role of external perturbations

4.5.1 Methods

The goal of Experiment 4 is to investigate the role of external perturbations during the learning stage (i.e., during the SMTs) and validate the hypothesis that an environment free of external perturbations (such as that provided by the uterus) provides favorable conditions for the development of the appropriate reflex circuitry (see Sect. 2.3). For this purpose, we identify the reflex connectivity as in Experiment 1 but we introduce small perturbations in the system during the learning process. These perturbations consist of short contractions with constant intensity ($M = 1$ mu) carried out by the external muscles, $M_{\text{Ext.1}}$ and $M_{\text{Ext.2}}$, which are activated at any given time step with a given probability. We varied the perturbation probability from 0 to 1 at 0.05 intervals.

4.5.2 Results

Our results indicate that for perturbation probabilities higher than 0.1, we cannot obtain the same reflex circuitry as in Experiment 1. This is because external perturbations induce information in the sensor signals which interfere with those produced by the SMTs. Figure 13 shows the connectivity obtained for perturbations generated with a probability of 0.2. As can be seen connections are identified between length sensors in the right leg and muscles in the left leg. These connections are incorrect as no transfer of information occurs between the two legs. In addition, the system identifies connections with length sensors which are not grounded in the physical system (e.g., the excitatory connection between M_{LG} and $\dot{S}^L_{\text{LI}}$). The force connectivity is similar to that identified in Experiment 1 because changes in force are detected only when the muscle is contracted.

4.6 Experiment 5: the reflex connectivity using different muscle models

4.6.1 Methods

The goal of Experiment 5 is to investigate the role of the muscle morphology on the reflex connectivity. For this purpose, the system learns the reflex circuitry as in Experiment 1 but using the artificial muscle model instead of the Hill-type muscle model (see Sect. 3.1).

4.6.2 Results

Figure 14 shows the connectivity obtained using the artificial muscle model. The connectivity obtained between the length sensors is the same as in the human muscle model. This result is expected as when one muscle contracts it extends the elastic shock cord of the antagonist muscle while decreasing the length of the homonymous muscle. On the force sensor, the connectivity observed includes extra inhibitory connections with the antagonist muscles. This occurs because when contracting one muscle the tension on the antagonist muscle also increases due to the lack of back-drivability in the motors. This experimental result suggests that the use of Hill-type muscles would be essential to self-organize the reflexes we observe in biology; muscles with different mechanical properties would lead to a different neural circuitry.

5 Discussion

5.1 Reflex connectivity

The main goal of this article is to show that circuitry of three basic reflexes, in two different sensor modalities, can be autonomously developed using SMTs and a simple correlation-based learning method. The success of our framework in self-organizing meaningful reflex circuitry provides a double contribution. On the one hand, it provides a clear testable hypothesis for the development of these reflexes in natural

systems; one relatively simple experiment would be to impair the normal functioning of the Ia and Ib at early stages of development and to investigate the subsequent development of the reflex circuitry. On the other hand, our framework provides a mechanism that can automatically endow artificial systems with useful feedback responses.

5.2 Biological plausibility

While our approach so far is overly simplified, we gained a number of insights into the motor development of biological systems. First, our method is based on a well-established plasticity rule, which is a variant of Hebbian learning similar to that used in the MDSI framework (Petersson et al. 2003). However, any variant of Hebbian learning is essentially a rate-based mechanism in which connections are allowed to change their nature (i.e., from excitatory to inhibitory or vice versa) simply by changing their signal. In biology, the nature of a connection is pre-established and it cannot be modified. This might be of fundamental importance for us as in our framework the nature of each reflex connection is established by our anti-Hebbian rule. We believe that establishing an appropriate (and biologically grounded) network structure as well as a plausible learning rule (e.g., a variant of spike-timing dependent plasticity) are the greatest future challenges of our approach.

Second, the sensor and motor models we use are only approximations of their biological counterparts. With the necessary augmentation, and with more precise biological models, we expect to learn more details about the degree to which this approach could hold in biological systems.

5.3 Timing, thresholding, linearity, and modulation of reflex activity

Although this article focuses mainly on the developmental process of the reflex circuitry, we need to discuss further the resultant behaviors observed. First, different reflexes appear at different times depending on the number of interneurons they entail. For example, the reverse myotatic reflex takes longer to be active than the myotatic reflex because it entails one interneuron in its pathway while the myotatic entails none. In our simulation, the reverse myotatic reflex also appears after the myotatic reflex, but this is because the force sensors are active only after the muscles start contracting (although this is also valid in biology).

Second, it is quite unlikely that the motor signal is a linear combination of the sensor values (see Eq. 9). Typically, a nonlinear function such as the sigmoid is used. Although that can easily be incorporated into our system, we see no direct benefit in doing so at such a preliminary stage of our investigation.

Third, and most importantly, all the reflex connectivity seems to be modulated by the supra-spinal systems according to the behavior being executed (Gottlieb and Agarwal 1973; Stein and Thompson 2006), i.e., the weights of all the connections can be manipulated from hierarchically superior systems (Hultborn 2001). For example, it is well known that the activity of γ-motoneurons can modulate the intensity of afferent Ia fibers. This is important because it reduces the relevance of the exact strength of the connectivity identified, and places a highest emphasis on the nature of the connectivity identified (inhibitory or excitatory) which in biological systems cannot be modified.

5.4 Emergence of reflex activity using the artificial muscle model

In the experiment carried out using the artificial muscle model, we can see that our approach scales rather well to produce relevant reflex activity in the artificial system. In contrast to the human muscles, the artificial muscles require an active extension. In this context, negative motor signals will cause the motors to actively drive backwards, allowing for an analogous of a muscle extension. Although the connectivity obtained with the artificial muscle model is different from that of the biologically inspired model, the behavior expected is similar (see Fig. 14). Relative to the myotatic reflex, muscle stretches caused by external loads will produce a contraction of the homonymous muscle. Relative to the reciprocal inhibition reflex, muscle stretches caused by external loads will produce an active extension of the antagonist muscle. Relative to the reciprocal inhibition reflex, increased muscle tensions will relax actively both the agonist and the antagonist muscles which results in a decrease of the overall tension in both muscles. This suggests that our framework can establish the appropriate reflex circuitry for a given body morphology (see also Sect. 5.6).

5.5 Dimensionality reduction

As argued in Földiák 1990, the anti-Hebbian learning rule has a potential for dimensionality reduction, by creating sparse sensorimotor circuitry. Our results offer support to this claim. If one analyses the connectivity matrix resulting from Experiment 1 (Fig. 9b) or Experiment 3 (Fig. 12), it is clear that the sensorimotor connectivity obtained is restricted to elements which interact at the local level, i.e., there is no sensor from one leg that has a connection to a muscle in the other. In the case of Fig. 9, the algorithm is able to transform a system which initially has 32 possible connections (4 motors $\times$ 8 sensors) into two separate systems with a total of 12 connections; in the case of Fig. 12, the total number of connections has been reduced from 128 possible connections to 40 con-

nections. Our experiments support the notion that Hebbian learning can be used to identify and segregate muscles in distinct groups of interacting homonymous, synergistic, and antagonist muscles.

5.6 Homeostasis and negative feedback reflexes

Our results suggest that our framework is particularly suited to produce homeostatic responses, or negative feedback reflexes. The negative signal entailed by the anti-Hebbian learning rule allows a system, natural or artificial, to learn the muscles which can counteract a given sensory stimulus, and thus maintain homeostasis. In this way, we believe that the exact same principle can be used to learn similar reflexes in other sensor modalities, such as pupillary light reflex in vision, or the acoustic reflex in audition.

6 Conclusions and outlook

The main hypothesis of this article was that three basic spinal reflexes could be self-organized from the interaction among five biologically inspired models: a musculoskeletal model (and its environment), a peripheral (sensory and motor) model, a model of SMTs, a learning model based on the correlations between sensor and motor activities, and a model of the reflex connectivity. We have shown that the circuitry of the reflexes obtained is consistent with that observed in biological systems, and that the connectivity obtained depends on the environment as well as on the morphology of the system. In addition, we addressed the potential of our framework to perform dimensionality reduction on the control signals, and thus have a potential application in the domain of artificial systems.

We are currently following two main research directions: one scientific and the other technological. The first is to make all the components of our system as biologically plausible as possible to show that in principle spinal reflexes can be self-organized in nature. The second is to analyze the functional significance of reflexes for voluntary movement and thus to contribute to the development of more robust artificial systems.

Acknowledgments The research leading to these results has received funding from the NCCR Robotics, as well as from the European Community's Seventh Framework Programme FP7/2007–2013—Challenge 2 —Cognitive Systems, Interaction, Robotics—-under Grant Agreement No. 231864—ECCEROBOT—and Grant Agreement No. 207212—eSMCs. We would also like to thank Matej Hoffmann and Naveen Kuppuswamy for reviewing the manuscript, and the reviewers of this article who effectively contributed to the results presented.

References

Bear M, Connors B, Paradiso M (2007) Neuroscience, 3rd edn. Lippincott Williams and Wilkins, London

Berniker M, Jarc A, Bizzi E, Tresch M (2009) Simplified and effective motor control based on muscle synergies to exploit musculoskeletal dynamics. Proc Natl Acad Sci USA 106:7601–7606

Bernstein N (1967) The co-ordination and regulation of movements. Pergamon Press, Oxford

Blumberg MS, Lucas DE (1994) Dual mechanisms of twitching during sleep in neonatal rats. Behav Neurosci 108:1196–1202

Blumberg MS, Lucas DE (1996) A developmental and component analysis of active sleep. Dev Psychobiol 1:1–22

Chen HH, Hippenmeyer S, Arber S, Frank E (2003) Development of the monosynaptic stretch reflex circuit. Curr Opin Neurobiol 13:96–102

Földiák P (1990) Forming sparse representations by local anti-hebbian learning. Biol Cybern 64:165–170

Gottlieb G, Agarwal G (1973) Modulation of postural reflexes by voluntary movement. J Neurol Neurosurg Psychiatr 36:529–539

Grillner S (2004) Muscle twitches during sleep shape the precise modules of the withdrawal reflex. Trends Neurosci 27(4):169–171

Hadders-Algra M (2004) Putative neural substrate of normal and abnormal general movements. J Pediatr 145:S12–S16

Hadders-Algra M (2007) Putative neural substrate of normal and abnormal general movements. Neurosci Biobehav Rev 31:1180–1190

Hebb D (1949) The organization of behavior. Wiley, New York

Hill AV (1938) The heat of shortening and dynamics constants of muscles. Proc R Soc Lond B 126(843):136–195

Holland O, Knight R (2006) The anthropomimetic principle. In: Proceedings of the AISB06 symposium on biologically inspired robotics

Hulliger M (1984) The mammalian muscle spindle and its central control. Rev Physiol Biochem Pharmacol 101:1–110

Hultborn H (2001) State-dependent modulation of sensory feedback. J Physiol 533(1):5–13

Jami L (1992) Golgi tendon organs in mammalian skeletal muscle: functional properties and central actions. Physiol Rev 72(3):632–666

Kosko B (1986) Differential hebbian learning. AIP Conference Proceedings, vol 151. pp 277–282

Kudo N, Yamada T (1985) Development of the monosynaptic stretch reflex in the rat: an in vitro study. J Physiol 369:127–144

Kuniyoshi Y, Sangawa S (2006) Early motor development from partially ordered neural-body dynamics: experiments with as cortico-spinal-musculo-skeletal. Biol Cybern 95:589–605

Latash ML (2008) Evolution of motor control: from reflexes and motor programs to the equilibrium-point hypothesis. J Hum Kinet 19: 3–24

Levinsson A, Garwicz M, Schouenborg J (1999) Sensorimotor transformation in cat nociceptive withdrawal reflex system. Eur J Neurosci 11:4327–4332

Lungarella M, Sporns O (2005) Methods for quantifying the information structure of sensory and motor data. Neuroinformatics 3(3):243–262

Mori H, Kuniyoshi Y (2010) A human fetus development simulation: self-organization of behaviours through tactile sensation. In: IEEE 9th international conference on development and learning. pp 82–87

Myklebust B, Gottlieb G (1993) Development of the stretch reflex in the newborn: reciprocal excitation and reflex irradiation. Child Dev 64(4):1036–1045

Petersson P, Waldenström A, Fåhraeus C, Schouenborg J (2003) Spontaneous muscle twitches during sleep guide spinal self-organization. Nature 424:72–75

Pierrot-Deseilligny E, Burke D (2005) The circuitry of the human spinal cord. Cambridge University Press, Cambridge

Prechtl HF (1990) Qualitative changes of spontaneous movements in fetus and preterm infant are a marker of neurological dysfunction. Early Hum Dev 23:151–158

Robinson SR, Brumley M (2005) Prenatal behavior. In: Whishaw I, Kolb B (eds) The behaviour of the laboratory rat: a handbook with tests, chap 24. Oxford University Press, Oxford, pp 257–65

Robinson SR, Blumberg MS, Lane MS, Kreber LA (2000) Spontaneous motor activity in fetal and infant rats is organized into discrete multilimb bouts. Behav Neurosci 114(2):328–336

Schouenborg J, Weng HR (1994) Sensorimotor transformation in a spinal motor system. Exp Brain Res 100:170–174

Sharma J, Angelucci A, Sur M (2000) Induction of visual orientation modules in auditory cortex. Nature 404:841–847

Smotherman W, Robinson S (1992) Prenatal experience with milk: fetal behavior and endogenous opioid systems. Neurosci Biobehav Rev 16:351–364

Smotherman W, Robinson S (1996) The development of behaviour before birth. Dev Psychol 32(3):425–434

Smotherman W, Robinson S (1998) Prenatal ontogeny of sensory responsiveness and learning. In: Greenberg G, Haraway M (eds) Comparative psychology: a handbook. Garland Publishing, Inc., New York, pp 586–601

Sporns O, Edelman G (1993) Solving bernstein's problem: a proposal for the development of coordinated movement by selection. Child Dev 64:960–981

Stein R, Thompson A (2006) Muscle reflexes and motion: how, what, and why. Exerc Sport Sci Rev 34(4):145–153

Todorov E, Ghahrammani Z (2003) Unsupervised learning of sensory-motor primitives. In: Proceedings of the 25th annual international conference of the IEEE engineering in medicine and biology society

Wang X, Merzenich M, Sameshima K, Jenkins W (1995) Remodelling of hand representation in adult cortex determined by timing of tactile stimulation. Nature 378:71–75

Wittmeier S, Jäntsch M, Dalamagkidis K, Knoll A (2011a) Physics-based modeling of an anthropomimetic robot. In: IEEE/RSJ international conference on intelligent robots and systems

Wittmeier S, Michael J, Dalamagkidis K, Rickert M, Marques HG, Knoll A (2011b) Caliper: a universal robot simulation framework for tendon-driven robots. In: IEEE/RSJ international conference on intelligent robots and systems

Zajac FE (1989) Muscle and tendon: properties, models, scaling and application to biomechanics and motor control. Crit Rev Biomed Eng 17(4):359–410