Eidgenössische Technische Hochschule Zürich
Swiss Federal Institute of Technology Zurich

Bio-Inspired Robotics Lab
Prof. Dr. Fumiya Iida

Semester-Project

Reflex-based Torso Stabilization

Spring Term 2013

Supervised by:
Dr. Hugo Gravato Marques

Author:
Christoph Bolliger

Contents

Abstract

In this thesis, the single-legged hopper discussed in [?] is extended with an upper body to explore the possibilities of combined torso stabilization and hopping. The original model used in [?] consists of a single leg and a hip which moves up and down on a vertical axis. The skeleton is driven by six muscles. In a first step, reflex connections (or weights) for the muscles are learned by cross-correlating the motor and sensor activity. This process is driven by spontaneous motor activity (SMA) which is a series of single muscle twitches. This has been implemented by twitching the muscles in a certain order and at regular intervals. Since there are two types of sensors in the muscles, type Ia (difference in length to resting length) and type II (rate of change in length), two weights will be learned. Once the weights have been learned, the feedback loop from the sensors to the muscles is closed through the weights. The behaviour of the model is modified by scaling those two weights with a pair of gains. For this hopper, a pair of gains was found that leads to stable hopping which was defined as hopping with an apex height of one meter at each hop.

The extending of the model happens by adding a rigid spine with an attached skull. To actuate the additional joint at the hip, a new pair of muscles is introduced. This extended model undergoes the same learning process as its predecessor. Parallel to the previously mentioned version of the SMA a more random one is tested to explore the influence of randomness on the learning process. In this new SMA the order in which the muscles are twitched is chosen randomly. This leads to a less smooth development of the weights during the learning process but yields quite similar weights. After the weights have been found, once again the feedback loop is closed. To find a pair of gains which both stabilizes the upper body and leads to stable hopping, the space of possible pairs of gains are searched in two different ways. One approach is a genetic algorithm, the other one a sweep search. The sweeps search produces a partial map where three areas can be distinguished. The majority of the search space leads to either spastic twitching or very weak responses to sensory input. A small area shows what I call standing behaviour which is quickly decaying hopping ending in a stable stance. The third area is an area where hopping is achieved, though I wasn't able to find a pair of gains for stable hopping. The torso stabilization is more forgiving than the hopping part and works with a large variety of gains.

Symbols

Symbols

w_{Ia}, w_{II}	weights for Type Ia and Type II sensors
G^{Ia}, G^{II}	corresponding gains for w_{Ia} and w_{II}
$S_{i,t}$	length of muscle i at time t
$\dot{S}_{i,t}$	rate of change of length of muscle i at time t

Acronyms and Abbreviations

ETH	Eidgenössische Technische Hochschule
BIRL	Bio Inspired Robots Lab
SMA	Spontanious Motor Activity

Chapter 1

Introduction

1.1 Motivation

Some time ago at Bio Inspired Robotics Lab (BIRL) a project for a bio-inspired reflex-based single-legged hopper has been started. Consisting of a hip on a verical axis and one leg actuated by six muscles, reflex connections have been learned by cross-correlating sensor and motor activity. The required muscle activity for this process was provided by spontaneous motor activity (SMA). This SMA was generated by twitching the muscles in a certain order and at regular intervals. Since there are two types of sensors in the muscles, type Ia (difference in length to resting length) and type II (rate of change in length), two sets of reflex connections (or weights) have been learned. At this point, the feedback loop form the sensors to the muscles were closed using the learned weights. By scaling the two weights with gains, stable 1D hopping has been achieved (see [?]), where stable hopping is defined as hopping with an apex height of one meter at each hop.
To take this model to 2D hopping (the hip moves freely), an upper body is required to enable the hopper to balance on the leg and keep upright. This thesis takes one step towards the 2D hopping by extending the model with a torso. With the hip still on a vertical axis, one pair of gains to stabilize the upper body and achieve stable hopping is looked for. The process remains the same as with the previous hopper. However, the learning process will be investigated more closely as well. A more random version of the SMA implementation will be tested where the order of the muscles is not fix any more. Taking the SMA verified in [?] as a reference, the influence of this added randomness on the learning process can be determined.

1.2 Hypothesis

The extension of the model as well as the learning of the weights with the SMA from [?] for the extended model are expected to work without any issues. The learning process with the new version of the SMAs is expected to work well, too.
Hopping and upper body stabilization should still be achievable with just one set of gains for the leg and the torso.

Chapter 2

Methods

The setup is illustrated in figure ??. In the following sections the implementations of the SMA Generator, the Musculoskeleton & Environment, the Reflex Circuits as well as the Supraspinal Systems are discussed.

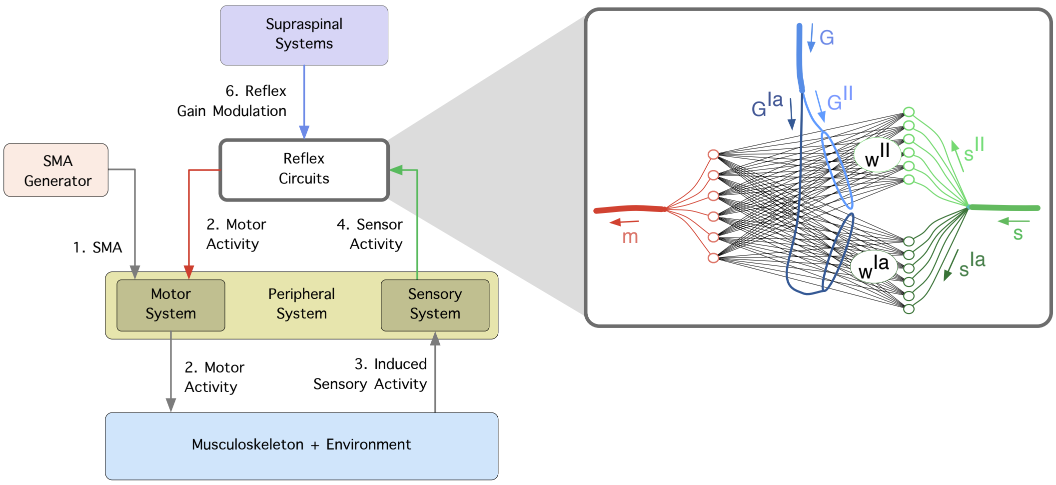

Figure 2.1: SMA (1) induce motor activity (2) which, through interaction with the Musculoskeleton and the Environment, induce Sensor Activity (3). To learn the Reflex Circuits, this sensor activity is cross-correlated with the motor activity to determine the weights (w^{IA} & w^{II}). Later, these learned reflex connections are scaled with their corresponding gains (G^{Ia} & G^{II}) to influence the behaviour of the musculosecletal system.(This schematics has kindly been provided by Hugo Gravato Marques and Arjun Bharadwaj)

2.1 SMA Generator

In previous work [?] with the hopper, the SMA is produced by the Pulse Generator Block of Matlab Simulink. For one iteration, all the muscles are twitched one after the other with an amplitude 0.01 and a duration of 0.3 second (see figure ??). The twitches take place at an interval of 5 seconds and the process is repeated forty times. Weights learned in this manner have been proofed to lead to stable hopping with the proper gains. The weights learned with this version of the SMA-generator will serve as a reference in this thesis.

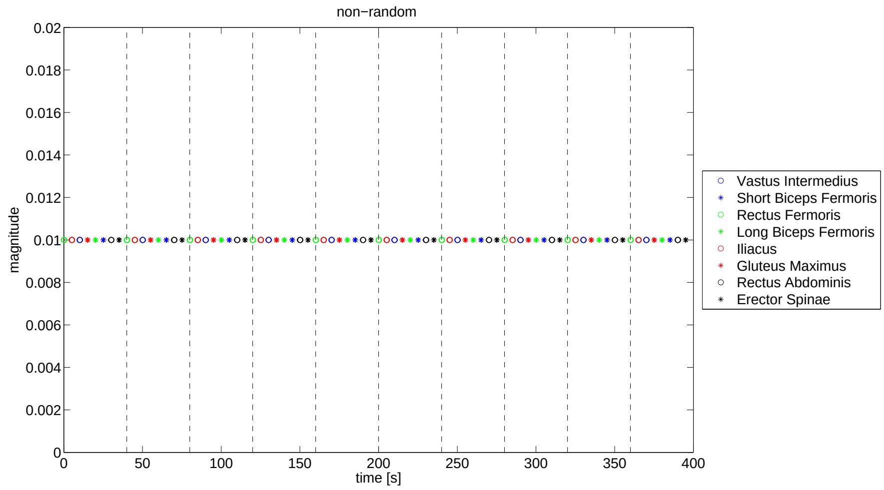

Figure 2.2: Nonrandom SMA: The dots represent the time a twitch takes place. Notice that every time segment (forty seconds) between two black lines contains the twitches of the eight muscles in the same order.

The SMA-signal is randomized by changing the order in which the muscles are twitched at each iteration. Magnitude and duration of the twitches as well as the interval between them remained unchanged (see fig.??). This way, a degree of randomness is introduced, but it is still ensured that each muscle will be twitched within a iteration. Discarding this restriction would greatly improve the randomness but there would be the possibility of one or more muscles not being twitched at all in a finite amount of time. If this was the case, the learned weights could not be used.

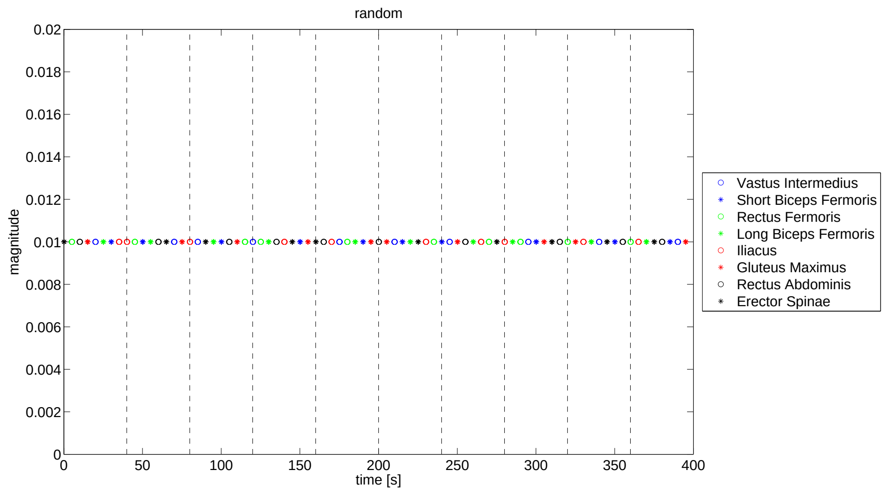

Figure 2.3: Randomized SMA: The order of the twitches within a time segment is now random.

2.2 Musculoskeleton & Environment

The model of the extended hopper consists of the head, torso and one leg which are actuated by eight muscles. It has been implemented in SimMechanics and possesses two ways of visualization. One is made up of simple geometries (see fig **??**), the other is based on CAD-models of human bones (see fig **??**). The second one slows down the simulations considerably, but it shows the muscles (white lines) and their activity (they get red).

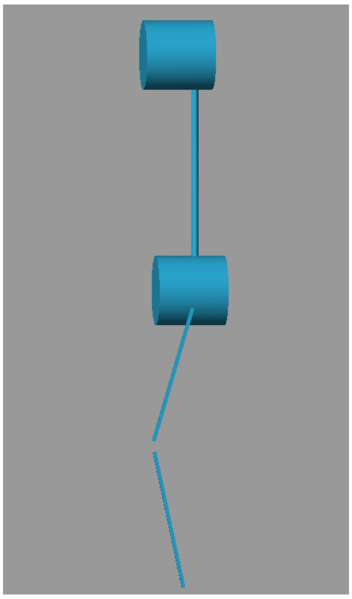

Figure 2.4: The Model in SimMechanics. It always runs during simulations and gives first impressions on promising candidates. Lacking information concerning the muscles, it is not fit for a more detailed review of a pair of gains.

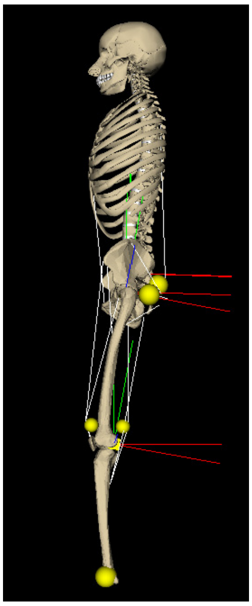

Figure 2.5: The Model in VR. The blue, red and green lines are the rigid bodies' coordinate systems and the white lines represent the muscles. Note that some of them lead over yellow spheres. They are pulleys to improve the angle of attack of some muscles.

2.2.1 The Skeleton

The skeleton is simplified to five rigid bodies (see figure **??**):

1 **Cranium**

2 **Spine** (the vertebrae and the ribcage are merged into one rigid body)

3 **Pelvis**

4 **Femur**

5 **Tibia**

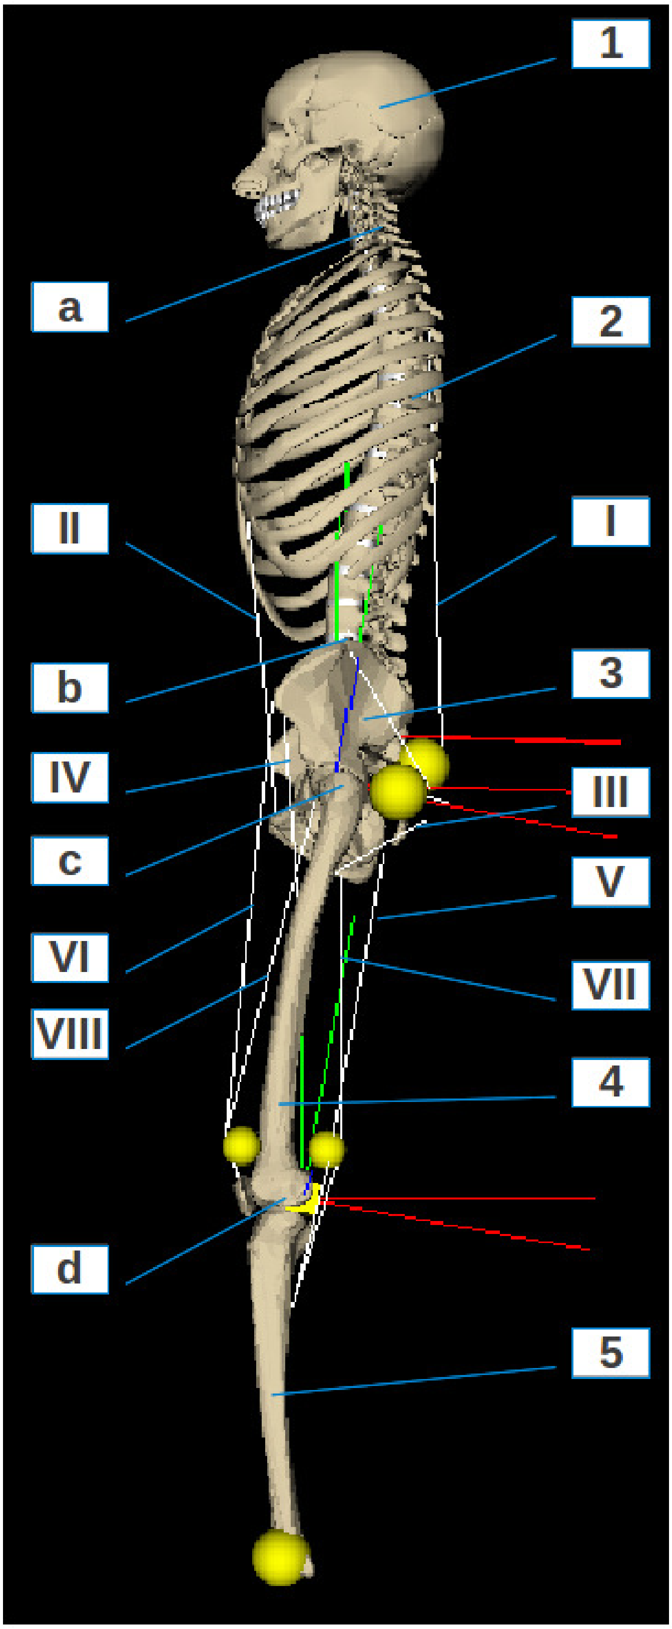

Figure 2.6: The Model in detail

All rigid bodies are assumed to be point masses of 0.5 kg, so inertias are neglected. They are connected by four joints (see figure **??**).

a Neck

b Spinebase

c Hip

d Knee

The neck is assumed to be stiff, all other joints are revolute joints. The angles of the joints are not restricted, since the muscles are expected to keep the joints within the natural operating range.

2.2.2 The Muscles

Each pair of muscles drives one or more joints:

- **Erector Spinae(I) & Rectus Abdominis(II)**
  Spinebase
- **Gluteus Maximus(III) & Iliacus(IV)**
  Hip
- **Long Biceps Femoris(V) & Rectus Femoris(VI)**
  Hip & Knee
- **Short Biceps Femoris(VII) & Vastus Intermedius(VIII)**
  Knee

The muscles are of the Hill-Type and deviation from resting length as well as rate of change of length are available as sensory data, mimicking the sensory feedback of type Ia and type II sensory fibres in real muscles.

2.2.3 The Environment

The musculoskeleton interacts in two ways with the environment. The pelvis moves on a vertical axis, preventing it from tipping in any direction or moving away from the spot and the lower end of the tibia interacts with the ground. To work on torso stabilization only, the hip may be fixed in space.

2.3 Reflex Circuit

For the learing process, the Anti-Oja rule (**??**) is used.

$$\Delta w_{ij,t} = -\eta \cdot m_{i,t} \cdot (s_{j,t} - m_{i,t} \cdot w_{ij,t}) \tag{2.1}$$

$w_{ij,t}$ is the reflex connection from sensor j to muscle i at time z. $m_{i,t}$ and $s_{j,t}$ are motor and sensor activity at time t of the corresponding muscles i and j respectively. The learning rate η has been set to 1000.
For learning the weights, the feedback loop is opened and the SMA-generator is used as input for the muscles. The output of the SMA-generator as well as the difference of the muscle length form the resting length or its derivative, depending whether the weights for the type Ia or type II muscle spindles are being learned, are fed into the Anti-Oja rule. The Anti-Oja rule updates the weights at every twitch and after a number of iterations, w^{Ia} or w^{II} will be learned. The development of the weights over time are logged, such that the learning process may be evaluated.

2.4 Supraspinal System

The reflexive reactions of the muscles to sensory input are calculated as described below.

$$m_{i,t} = G^{Ia} \cdot \sum_{j=1}^{N} s_{j,t}^{Ia} \cdot w_{ij}^{Ia} + G^{II} \cdot \sum_{j=1}^{N} s_{j,t}^{II} \cdot w_{ij}^{II} \tag{2.2}$$

$m_{i,t}$ is the motor activity of muscle i at time t. $s_{j,t}^{Ia}$ and $s_{j,t}^{II}$ are the sensor activity at time t of their respective muscle spindle type. w_{ij}^{Ia} and w_{ij}^{II} are the weights which define the reflex connections and G^{Ia} and G^{II} are the gains with which their respective feedback is scaled. Those two values represent the influence of the supraspinal system on the reflex circuit. With them, the exact behaviour of the model can be varied. To find the ones which lead to the desired behaviour (hopping with a stable torso), a huge space of possible pairs of gains needs to be considered. To make any kind of automated search possible, the results of the simulations with given gains have to be represented by a single number. To make this possible, the trajectory of the hip is analysed. The closer each pinnacle of a hop is to the desired hopping height, one meter, the better the hopping is. Based on that, the H-function was devised.

$$H = \frac{1}{N} \cdot \sum_{n=1}^{N} \Delta {h_n}^2 \tag{2.3}$$

or, more detailed,

$$H = \frac{1}{N} \cdot \sum_{n=1}^{N} (h_n - 1)^2 \tag{2.4}$$

N is the number of hops and h_n the height of the n^{th} hop.
The deviation from the desired hopping height is squared to eliminate the signs and the sum is normalized to prevent varying number of hops from distorting the results. With this H-function, a big amount of data can be evaluated at a glance.
Since every simulation takes quite long, a rule to detect a certain fail is established to save time. If the hip goes below 0.6 meters or above 2.5 meters, the simulation is aborted and a value of ten is assigned to the pair of gains. This way, fails are easily distinguishable from the rest of the results.
To ensure a stable upper body, no additional terms need to be added to the H-function, since an unstable upper body will swing down below the hip and collide with the leg, instantly ending any hopping.
To find the gains, two methods of automated search are considered, a genetic algorithm and a sweep search.

2.4.1 Genetic Algorithm

For the genetic Algorithm, a set of randomly chosen starting pairs are simulated and then evaluated with the H function (**??**). The pairs of gains are inherited to the next generation according to their H-value, the lowest values having the best chance to be passed on. After this, mutations (random changes to the values) take place to ensure diversity. To make sure that the most promising pair is not discarded by mutation, an elite is introduced.The elite, the best pair of an iteration, is passed on to the next generation without alterations.

2.4.2 Sweep Search

The sweep search is a simple brute force approach. A grid is laid over the search space and at every intersection a simulation is conducted. After each simulation, the data is stored and the algorithm proceeds to the next intersection. Since the size of the areas with the desired gain pairs in the search space is not known, a good resolution for the grid must be found experimentally.

Chapter 3

Results

3.1 The Weights

Using the nonrandom SMA, weights were learned for both the type Ia and the type II sensors. In figure **??** the development of Ia weights learned by twitching long biceps femoris with the nonrandom version of the SMA are plotted over time. Three types of weights are distinguishable. The positive ones of the agonists (other muscles that work with the one in question, e.g. short biceps femoris), the negative ones of antagonists (muscles that work against the one in question, e.g. rectus femoris) and the ones close to zero for muscles which only have a very weak one or no participation in the motion at all (e.g. rectus abdominis). The values converge within 400 seconds, which is equivalent to ten iterations.

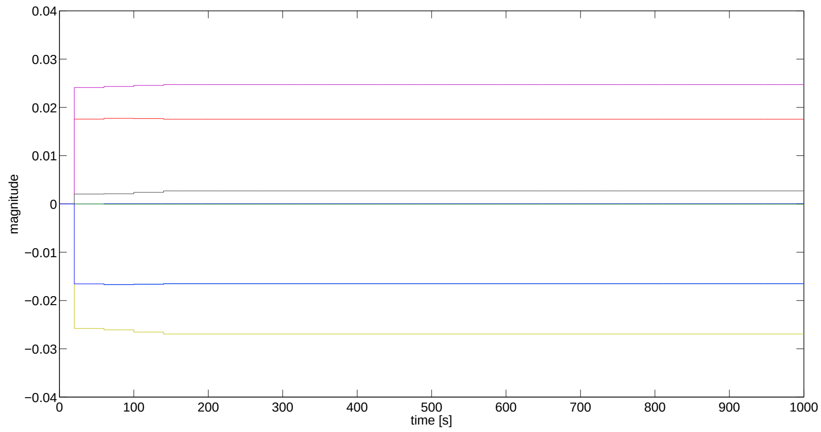

Figure 3.1: The weights over time. With the regular intervals between the twitches of the same muscle, the progress is quite smooth. Note the three types of connections: positive for agonists (pink & red), negative for antagonists (blue & light green) and weak connections (the rest).

When learning the weights with the random SMA, the process is less smooth, but similar behaviour can be observed (see figure **??**).

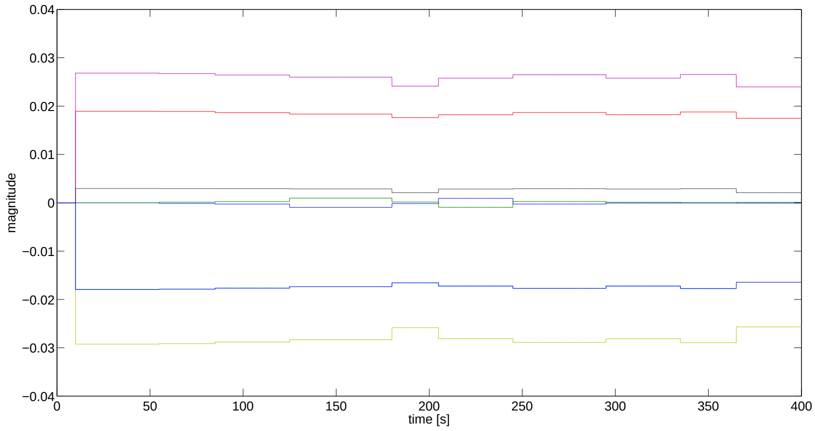

Figure 3.2: Weights over time. The development over time of the weights is less smooth than in figure ??

The two Ia weights from random and non-random SMA are plotted in figures ?? and ?? as Hinton-Matrices. The weights look similar to the naked eye but, after an overall analysis, they will be compared more closely.

For the Hinton-Matrix, the size of a circle represents the magnitude of the connection and the color the sign. White circles are positive values and grey are negative ones. All the diagonal elements have positive feedback.

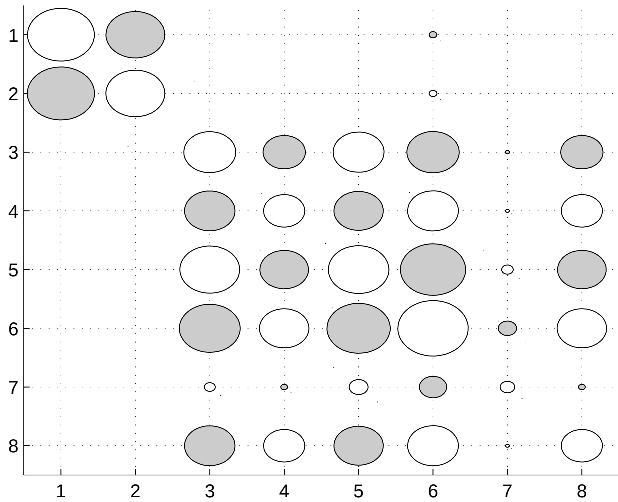

Figure 3.3: weights from nonrandom SMA

This is good, since negative values would mean that the muscles counteract themselves. Furthermore, there is a clear lack of connectivity between muscles one and two which actuate the torso and muscles three trough eight which actuate the leg.

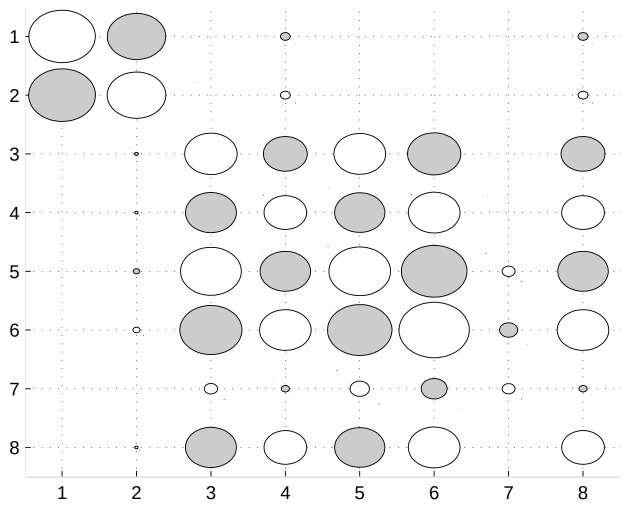

Figure 3.4: weights from random SMA

The pattern of diagonals altering colors from the middle outwards indicates that even numbered muscles have positive connections among each other. The same is true for uneven numbered muscles. The connections between the groups are all negative. This fits the model, since all the uneven numbered muscles are located at the back of the model and the even numbered ones at the front (see figure **??** on page **??**). The back ones are for moving torso and femur back and bending the knee while the front ones move the torso and the femur to the front and stretch the knee. For the weak connections of number seven (short biceps femoris) I can not offer an explanation. Over all, the weights seem reasonable and should make hopping as well as torso stabilization possible.

To compare the two sets of learned weights, a special version of the Hinton-Martix was created (see figure **??**). In this version, the size of a circle located at the intersection between two numbers corresponds to the difference between the weights for this sensor-muscle connection relative to the magnitude of the biggest weight value. The colors refer to same sign for white and opposite sign for grey. To enhance visibility of the colors, the data has been magnified by a factor of ten.

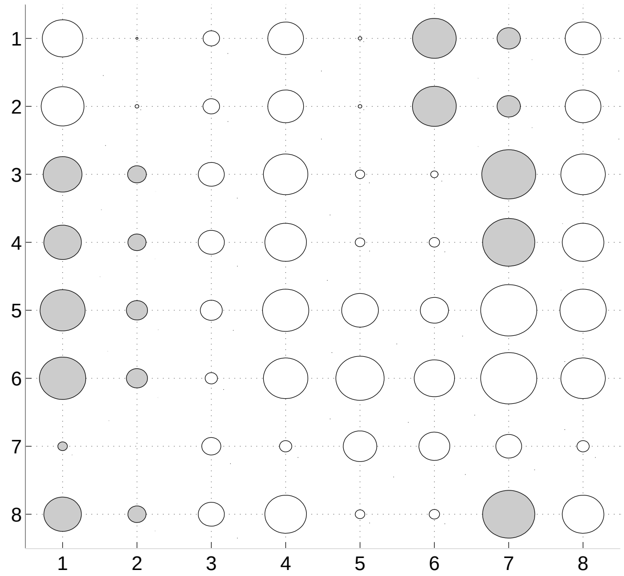

Figure 3.5: Comparison of the weights from random and nonradom SMA. Each muscle is referred to by the number assigned to it in section **??** on page **??**. White signifies same sign, grey opposite sign.

The biggest differences are located at the connections between the torso and the leg (intersections of one and two with three through eight), which is not relevant since those connections origin in noise (compare figure **??** & **??**). Further differences occur with number seven (short biceps femoris). The three most significant deviations (with opposite signs) occur at the weak connections (compare figures **??** and **??**).

3.2 The Gains & Hopping

3.2.1 Genetic Algorithm

For the genetic algorithm, the first generation of gain pairs was chosen randomly. At each iteration, a simulation is conducted and the hopping is evaluated by means of the H-function (see ??). Then the new generation is chosen out of the old one. Pairs of gains with a low H-value have a better chance to be passed on. After this, the values are modified to ensure diversity. After thirty iterations, the following data was found.

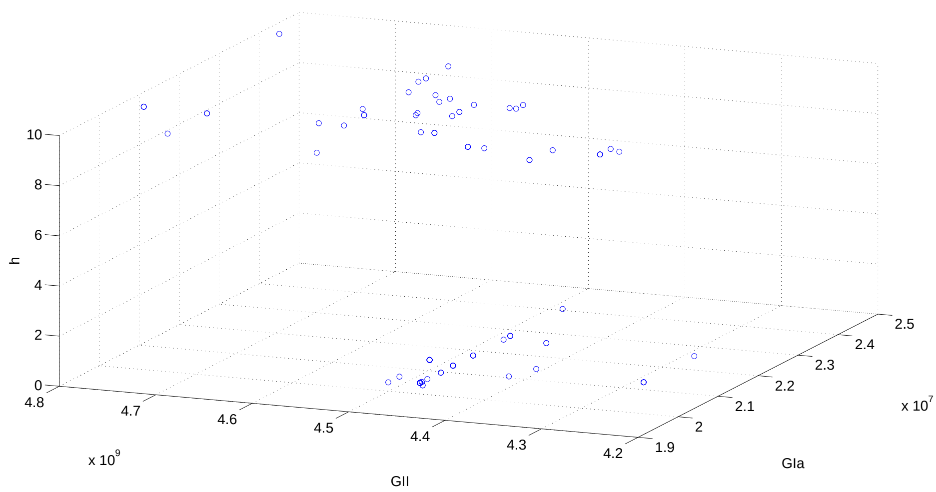

Figure 3.6: GA results: The algorithm settled into the area around $G^{Ia} = 2 \cdot 10^7$ and $G^{II} = 4.5 \cdot 10^9$

The algorithm settled into a local minima which, unfortunately, wasn't good enough for hopping. There seems to be a area of better values in a big space of failed simulations. The genetic algorithm may be suitable for the fine-tuning of the gains once a reasonably good set has been determined, but it is not very efficient at finding the good area in the whole search space.

3.2.2 Sweep Search

With a number of sweeps the space from $GIa = 0$ and $GII = 0$ to $GIa = 8.5 \cdot 10^7$ and $GII = 5.3 \cdot 10^9$ has been mostly mapped (see ??). Given the small values of H of the successful simulations (magnitude of 10^{-2}) the failed simulations H-values were scaled down from ten to 0.06 for better visualization.

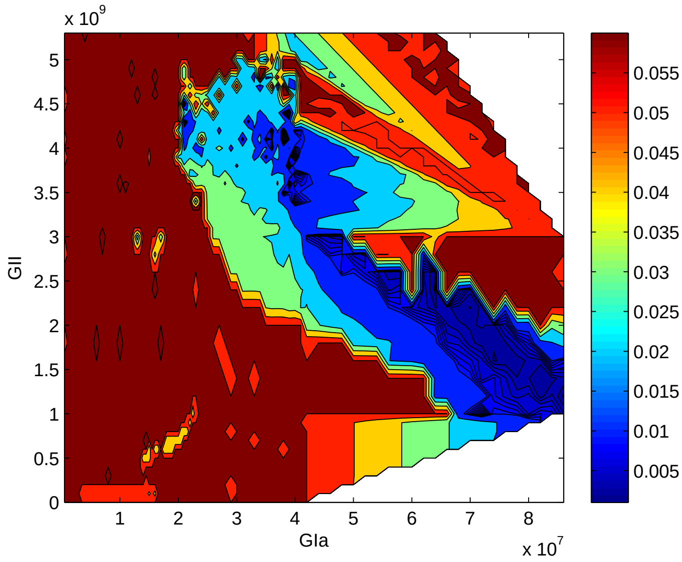

Figure 3.7: Map of the explored part of the gains space. The failed simulations have been scaled down to 0.06 for better visualization. Dark red values are failed simulations, the dark blue areas show hopping behaviour.

Big parts of the map (see figure **??**) are dark red (failed simulations). Besides of the big blue area to the right, the small yellow area in the lower left corner stirred my interest. This particular area is made up by gain pairs with similar G^{Ia} to G^{II} ratios.
A closer look (see figure **??**) shows that this area consists of values around 0.04 and has a G^{II} over G^{Ia} ratio of approximately 35. Analysis of the logged data showed that the pairs of gains in this area lead to a fast decaying hopping which ends in standing. The torso stabilizes very nicely in this setup.

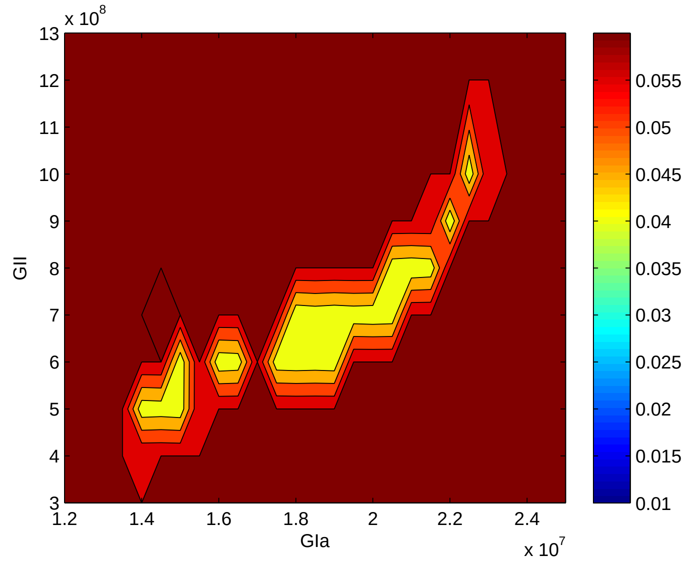

Figure 3.8: map of the hopping gains

The big blue area (see **??**) on the other hand has significantly better H-values. Those values lead to hopping, but by no means stable hopping. The hopping height varies and the duration of the hopping is below one minute. Here again, the torso stabilized nicely, but due to the continuing hopping, the equilibrium keeps being disturbed. The hip trajectory of the best set of gains can be found in figure **??**.

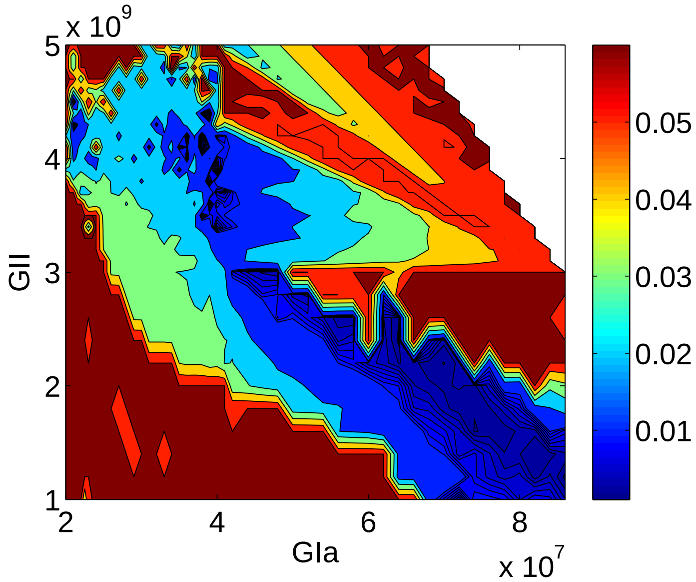

Figure 3.9: map of the hopping gains

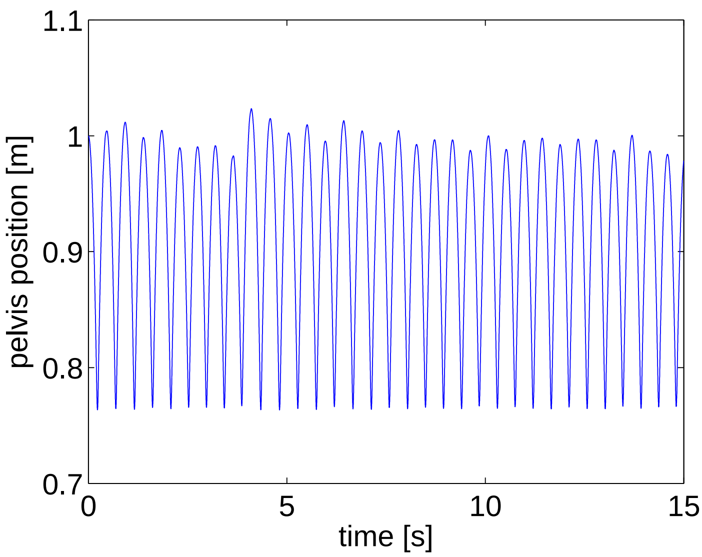

Figure 3.10: hip trajectory of $G^{Ia} = 21000000$ and $G^{II} = 4500000000$

Chapter 4

Conclusions

The fact that both standing and hopping, although the latter not stable, were achieved with the same weights proofs that different behaviours can be achieved by changing G_{Ia} and G_{II} only.
Even if the correct pair of gains for stable hopping couldn't be found within the given time-frame, I'm convinced that it exists. I see the areas with close to hopping gains as strong evidence and I believe the nature of the gains space should be further examined. A better understanding may speed up the finding of desired gains for future projects, but transferability of this to different models has to be examined as well.
The random order of the SMAs did not upset the learning of the weights, leaving room to go further towards more random SMAs in future projects.

Bibliography

[1] Arjun Bharadwaj : *From Spontaneous Motor Activity to hopping: a developmental scheme*. Master Thesis, ETH Zurich, July 2013.

[2] Hugo Gravato Marques, Kristin Völk, Stefan König and Fumiya Iida : *Self-Organization of Spinal Reflexes Involving Homonymous, Antagonist and Synergistic Interactions.*

[3] Hugo Gravato Marques, Farhan Imtiaz, Fumiya Iida, Rolf Pfeifer : *Self-organization of reflexive behavior from spontaneous motor activity*. Biological Cybernetics, DOI 10.1007/s00422-012-0521-7, 2012.