Berthouze, L., Kaplan, F., Kozima, H., Yano, H., Konczak, J., Metta, G., Nadel, J., Sandini, G., Stojanov, G. and Balkenius, C. (Eds.)
Proceedings of the Fifth International Workshop on Epigenetic Robotics: Modeling Cognitive Development in Robotic Systems
Lund University Cognitive Studies, 123. ISBN 91-974741-4-2

Developmental acquisition of entrainment skills in robot swinging using van der Pol oscillators

Paschalis Veskos
Department of Electrical and Electronic Engineering
Imperial College London SW7 2BT
paschalis.veskos@imperial.ac.uk

Yiannis Demiris
Department of Electrical and Electronic Engineering
Imperial College London SW7 2BT
y.demiris@imperial.ac.uk

Abstract

In this study we investigated the effects of different morphological configurations on a robot swinging task using van der Pol oscillators. The task was examined using two separate degrees of freedom (DoF), both in the presence and absence of neural entrainment. Neural entrainment stabilises the system, reduces time-to-steady state and relaxes the requirement for a strong coupling with the environment in order to achieve mechanical entrainment. It was found that staged release of the distal DoF does not have any benefits over using both DoF from the onset of the experimentation. On the contrary, it is *less* efficient, both with respect to the time needed to reach a stable oscillatory regime and the maximum amplitude it can achieve. The same neural architecture is successful in achieving neuromechanical entrainment for a robotic walking task.

1. Introduction

Non-linear differential equations have been used in robotics as central pattern generators (CPGs) for a multitude of tasks such as robot swinging (Lungarella and Berthouze 2002), arm motion (Williamson 1998) and locomotion (Taga 1991; Zielinska 1996; Lewis, Etienne-Cummings et al. 2003), both in simulation and robotic experiments. CPGs have become increasingly popular as they provide biologically-inspired, robust and adaptive motion.

With the use of CPGs in robot swinging, an interesting range of rhythmical behaviours can be explored. Thus, what might seem as an uncomplicated task offers the right amount of complexity to investigate and make comparisons with theoretical findings and different oscillators, while at the same time maintaining a controlled experimental environment.

An increasingly strong case has been made for embodiment and hence the use of real robots instead of simulations, initially with (Brooks 1991) and now many others such as (Thelen and Smith 1994; Beer, Chiel et al. 1998). The main argument is that the body, "which mediates perception and affects behaviour, plays an integral role in the emergence of human cognition, […] the central tenet of embodied cognition is that cognitive and behavioural processes emerge from the *reciprocal and dynamic coupling between the brain, body and environment*" (Lungarella, Metta et al. 2003). We endorse this school of thought and have thus made use of an embodied robotic platform to investigate our problem.

In a previous study (Lungarella and Berthouze 2002), the same task was investigated using a different non-linear oscillator (Matsuoka 1985). In this paper, we explore the behaviour of the system and its ability to achieve entrainment with the natural dynamics under different morphological configurations, with and without neural entrainment. The van der Pol oscillator (Strogatz 2001) was chosen for the smaller number of parameters requiring tuning, robustness (Matsuoka is a near-harmonic oscillator that does not feature an asymptotically stable limit cycle) and straightforward computational implementation.

2. The Degrees of Freedom Problem

The *degrees of freedom problem* was first attacked theoretically by Nikolai Bernstein in a series of papers, published collectively in *The coordination and regulation of movements* (Bernstein 1967). The main theme of his framework was the staged acquisition of control, in the proximal-to-distal direction. This is a three-staged process:

1. Distal (peripheral) degrees of freedom (DoF) are decreased (frozen) to a minimum.
2. Once control of the more proximal ones has been mastered, the frozen degrees of freedom are released and the process continues.
3. Once the organism can control its morphology in its entirety, it can begin to exploit any reactive phenomena that occur to its advantage.

This process is intuitive, e.g. the reduction of variables to be controlled is reasonably expected to accelerate acquisition of control and also draws a parallel to Piaget's notion of stages of cognitive development. As such, it has been widely accepted for a long time. There is some experimental verification of the process, with the most widely cited example being that of the disappearing infant stepping/kicking reflex. Exhibited by infants as young as 12 weeks *in utero*, this reflex disappears by the age of four or five months. It only reappears much later when the child learns how to walk at around 12 months. This is viewed as freezing of a distal DoF because the infant

cannot yet cope with a high number of variables and subsequent release due to attainment of control.

It was, however, later shown that the basis for this change in behaviour is not of neural, but of *morphological*, origin (Thelen and Smith 1994). That is, muscle strength is not able to cope with the dramatic leg mass increase accompanying infant maturation (Goldfield 1995). Furthermore, there is experimental evidence that Bernstein's 3-stage model is not universal. For some tasks, an additional stage of freezing and releasing some degrees of freedom may be required, depending on task constraints (Newell and Vaillancourt 2001).

Utilising the task of robot swinging (Lungarella and Berthouze 2002), we will investigate the effects of the distal-proximal DoFs on the task depending on their degree of coupling. We also aim to characterise the swinging task under different morphological configurations. For this purpose, we activate each degree of freedom separately and observe the resultant behaviour. Subsequently both hip and knee joints are activated, under different coupling schemes.

3. Experimental Setup

The robot is held from a horizontal bar like a pendulum so that it can swing freely about it. A coloured marker is placed on the robot so that a webcam viewing the setup from the side can track the marker's position. The x coordinate of this marker is then used as feedback for the neural oscillator. In this study, only the hip and knee joints were actuated, while all others on the robot were held stiff. This way the system can be viewed as an underactuated triple pendulum with the top joint being free while the bottom two joints are totally forced to the output of the nonlinear oscillator. The experimental setup and an equivalent representation are shown in Figure 1.

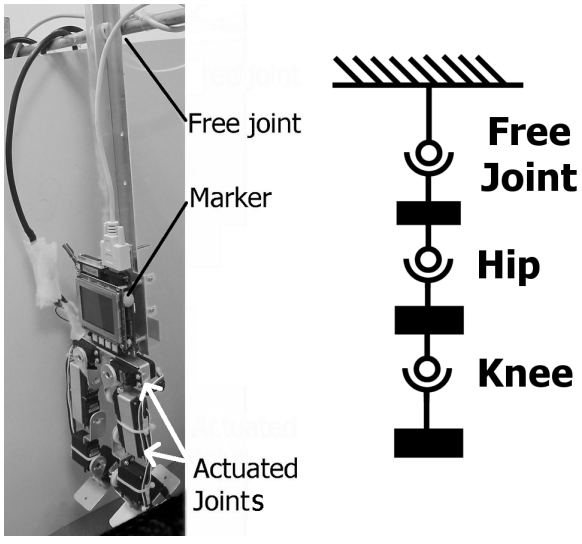

Figure 1: The experimental setup and equivalent representation.

It should be noted that this equivalent representation is only shown here for clarity and was not used as a model of the system. This is in accordance with a major design principle in the ontogenetic paradigm: The explicit avoidance of directly engineered approaches, as these stem from the designer's understanding of robot physics and task dynamics. Instead of hard-wiring the problem's solution in the robot's brain, the "[control] structure should reflect the robot's own process of understanding through interactions with the environment" (Asada, MacDorman et al. 2001).

4. Non-linear Oscillator

The equations of the van der Pol oscillator, as used in our experiments, are:

$$\ddot{x}_{hip} + \mu\left(x_{hip}^2 - 1\right)\cdot \dot{x}_{hip} + \omega^2 x_{hip} = G_{in}\cdot fb + G_{hip-knee}\cdot x_{knee} \quad (1)$$

$$\ddot{x}_{knee} + \mu\left(x_{knee}^2 - 1\right)\cdot \dot{x}_{knee} + \omega^2 x_{knee} = G_{in}\cdot fb + G_{knee-hip}\cdot x_{hip} \quad (2)$$

where $\mu \geq 0$ is a parameter controlling the damping term, ω is the natural frequency of the oscillator, fb represents the feedback from the vision system, G_{in} is the feedback gain, while $G_{hip-knee}$ and $G_{knee-hip}$ are the cross-coupling term gains. The final output given to the position-controlled motors activating the joints, is:

$$\theta_i = G_{out}\cdot sign\left(\dot{x}_i\right),\ i = \{hip, knee\} \quad (3)$$

where G_{out} is the output gain.

Using the derivative of the equation's solution to generate the motor command has the effect of removing any DC components from the oscillator signal; taking the sign of this result generates a binary command for full joint extension or contraction (Figure 2).

This is necessary for the low bandwidth motors that we use and also avoids the need for signal normalisation: unlike other nonlinear oscillators, the van der Pol's output varies in amplitude, dependant on the amplitude of its inputs (Williamson 1999). We thus convert the periodic oscillator signal to a pulse-width modulated square wave retaining frequency information.

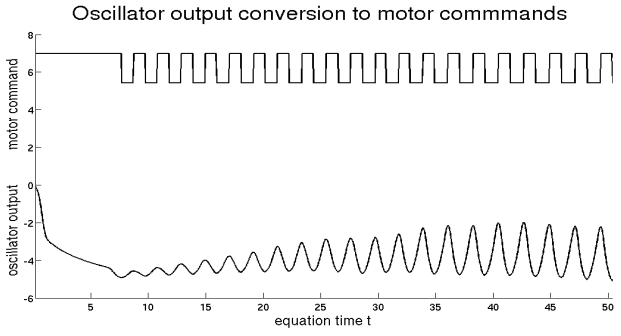

Figure 2: Example oscillator output (e.g. xhip, bottom waveform) and corresponding motor commands (θ, top waveform). Note that the drift in the DC component of the oscillator output does not affect the motor command.

By altering the gains, a range of different behaviours can be realised. If the value of G_{in} is too low, then the oscillator does not entrain to the feedback signal. Likewise, G_{out} has to be sufficiently high to 'excite' the

mechanical system, yet given the low actuator bandwidth not so high as to make the distance the joints have to move excessively large. The gains affecting the cross-coupling terms control the influence each oscillator has on its counterpart. If they are not equal to each other, a phase difference between the two oscillators is introduced. To avoid this suboptimal behaviour, they were kept equal throughout this study.

Unless mentioned otherwise, the following values were used throughout the study: $\mu = 1.00$, $\omega^2 = 1.00$, $G_{in} = 0.143$, $G_{hip-knee} = G_{knee-hip} = 0.500$ and $G_{out} = 0.784$. The initial conditions given to the numerical integrators throughout this study were: $\{x_{hip}, \dot{x}_{hip}, x_{knee}, \dot{x}_{knee},\} = \{0.00, 0.00, 0.00, 0.00\}$.

5. 1-DoF experimentation

Initially the system was only allowed to actuate the hip joint. As shown in Figure 3, after a transient phase of increasing amplitude oscillations, the system settles into a stable oscillatory regime. The maximum amplitude reached was 163 units.

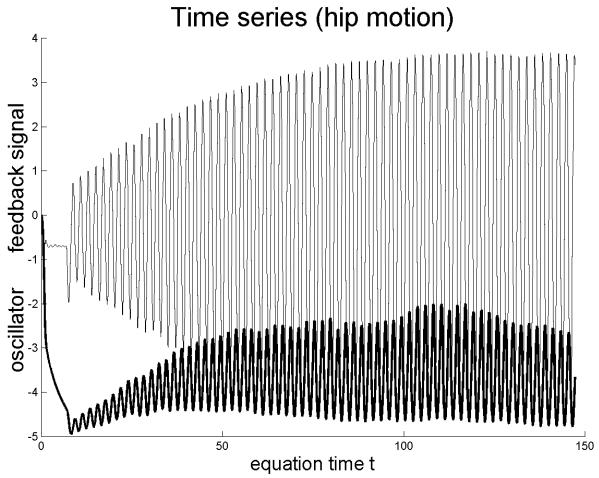

Figure 3: The time series with the hip actuated. The top signal is fb and the bottom signal is oscillator output (x_{hip}).

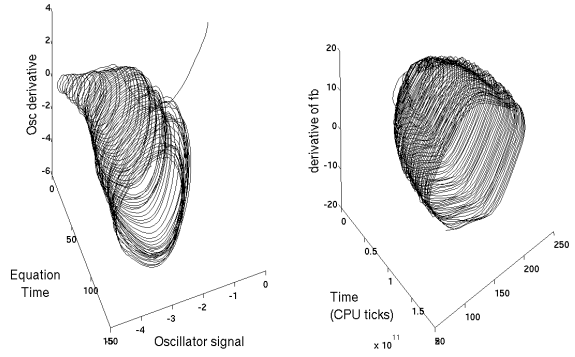

Figure 4: Phase plots for the oscillator (left) and feedback signals (right) for a typical experimental run where the hip joint was actuated.

Next, the system was run under the same setup, but with the output of the oscillator fed to the knee joints. Due to geometry and the fact that the motors move a smaller mass, the resultant oscillations were much smaller in amplitude Figure 5.

This way the feedback gain had to be significantly increased to 0.50 for entrainment to take place. The maximum oscillation amplitude attained in this setup was just 61 units. The mechanical system limit cycle contains multiple closed circular regions per iteration (knots), suggesting suboptimal behaviour.

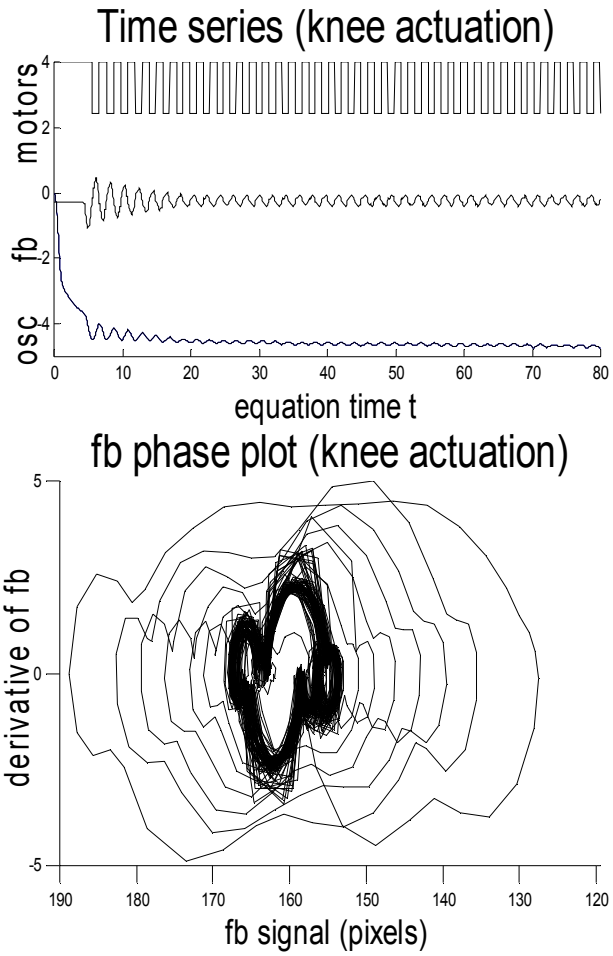

Figure 5: Actuating the knee joints alone resulted in small amplitude oscillations. On the left, the time series is shown; from top to bottom, the signals are: knee servo command (θ), mechanical system position (fb) and oscillator output (xknee). On the right is a 2D projection of the fb limit cycle. The isolated external trajectories belong to the initial oscillations

6. Independent 2-DoF experiments

Subsequently, both joints were actuated at the same time. It was found that the system was able to perform much better in the 2-DoF configurations than the 1-DoF ones. Oscillations were more stable, of larger amplitude and the steady state was attained faster.

With the cross-coupling constants $G_{hip\text{-}knee}$ and $G_{knee\text{-}hip}$ set to 0, the two oscillators are only linked by means of the mechanical feedback signal. It was thus necessary to set the feedback gain parameter, $G_{in,}$, to a relatively high value (0.5) to allow sensory feedback to modulate the oscillator's output. This results in a stable oscillatory

regime that attains maximal amplitude oscillations of 201 units (Figure 6).

The regime is, however, prone to sporadic glitches during the transient phase, indicating that it is not totally stable. The time series and corresponding limit cycle for such a trial can be seen in Figure 6. The glitches manifest themselves as single out-of-phase kicks, in addition to normal behaviour.

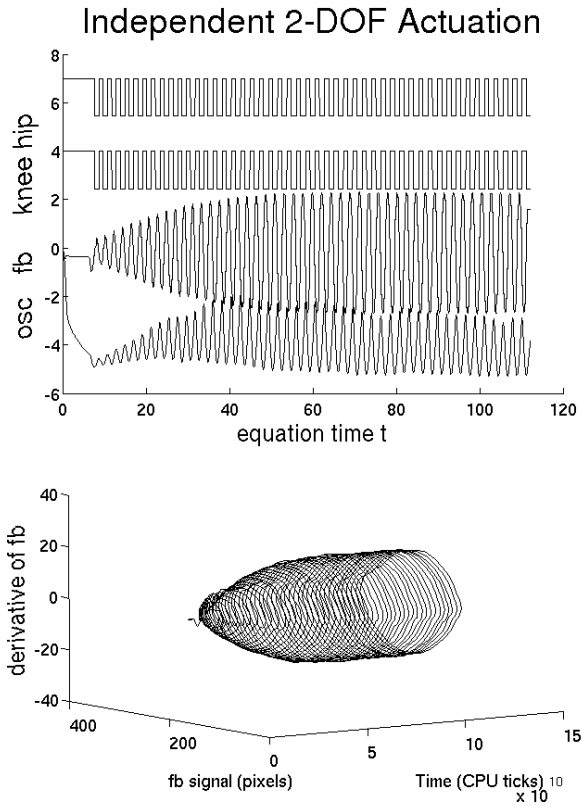

Figure 6: Time series (left, top to bottom the signals shown are knee- and hip- servo commands, feedback signal and oscillator output) and mechanical phase plot (right) of the stable oscillatory regime obtained with cross-coupling *disabled*.

7. Neural entrainment

With the cross-coupling gains both set to 0.50, an explicit connection is made between the two oscillators. The output of each one is fed directly into the other, before exciting the motors. With this connection, the system finds a very stable oscillatory regime, even for low feedback gain values. Furthermore, the system reaches its steady-state faster and attains the largest maximal amplitude (206 units) among all experimental configurations. The time series and mechanical system phase plot can be seen in Figure 7.

The corresponding phase plots for this experiment can be seen in Figure 8. The mechanical limit cycle is now smoother and varies less with time, but the neural system exhibits a low-frequency oscillation enveloping the normal behaviour.

Provided that the feedback gain is high enough, the system can always reach a stable regime. The presence of neural entrainment acts as a stabiliser, removing glitches in the oscillator output and avoiding sudden transient behaviour. With the cross-coupling gains set to 0.25 the system achieves a maximum oscillation amplitude of 198 units and with the gains set to 0.10, 200 units.

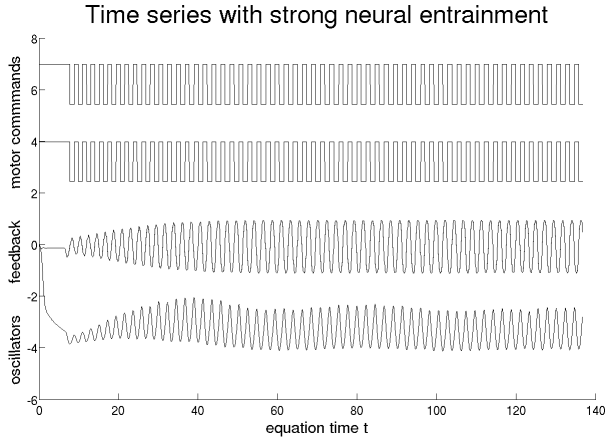

Figure 7: Time series of the stable oscillatory regime obtained with cross-coupling *enabled*. Top to bottom the signals shown are knee- and hip- servo commands, feedback signal and oscillator output.

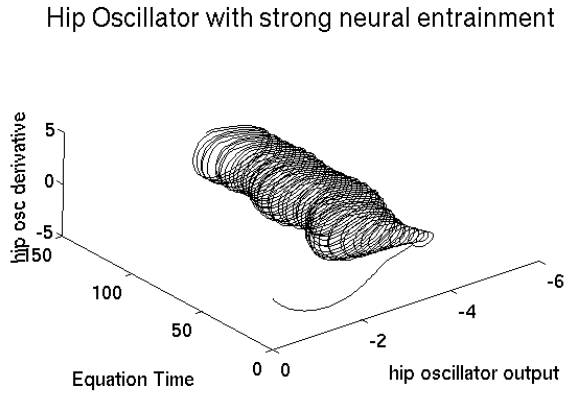

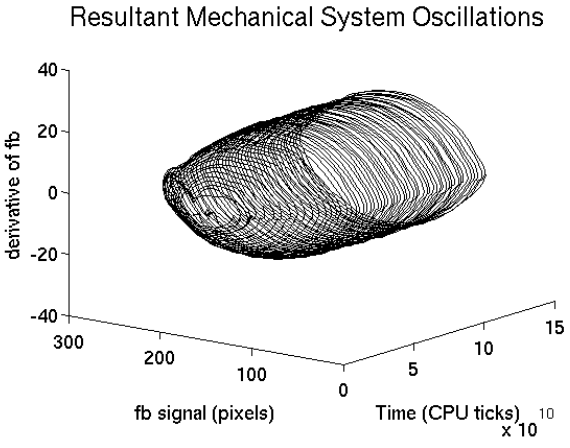

Figure 8: Phase plots for the oscillator output (top) and mechanical feedback (bottom) signals in the presence of strong neural entrainment (cross-coupling gains both set to 0.50).

8. Staged distal DoF release

Next, a staged release of the knee DoF was carried out. The system initially started oscillating using only the hip, with the knee output and feedback input inhibited. In two

different trials, the oscillator was enabled once the system had reached its steady state and at an earlier stage.

The knee oscillator is very quick to entrain to the feedback signal, within one (neural system) period; this is expected as mechanical entrainment has already been achieved. The presence of neural entrainment improves upon this only marginally. The staged release of the second DoF has the effect of transitioning the system to a slightly suboptimal version of the 2-DoF oscillatory mode. This configuration achieves both a smaller maximum amplitude (by about 5%, 190 units) and requires far more time to reach the steady state.

An aggregate performance comparison for all different morphological configurations can be seen in Figure 9. The best performance is achieved by using both DoF and strong neural entrainment, closely followed by the independent- and the two staged, 2-DoF configurations. The single-DoF setups exhibit the worst performance.

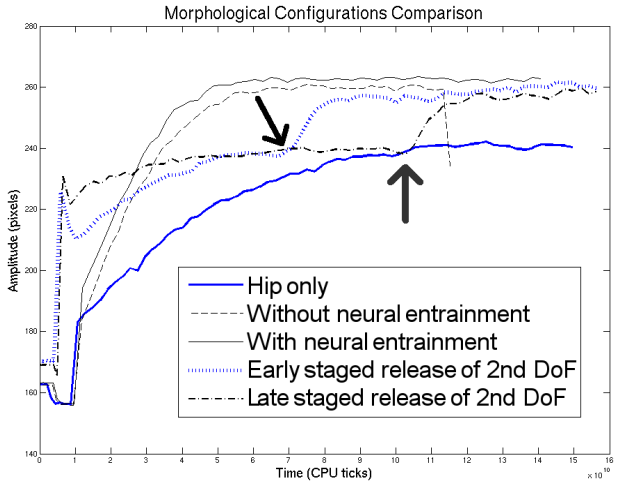

Figure 9: Performance comparison for the 5 morphological configurations. Starting the experimental runs with both DoF enabled offers the best performance, both in terms of maximum oscillation amplitude attained and time to reach the steady state. The arrows denote the instant where the second DoF was released and its oscillator activated.

9. Towards acquiring walking skills

The same neural architecture was also used in a robotic walking task. The experimental setup can be seen in Figure 10. The robot is supported by a metal arm and thus restricted to move in the sagittal plane, i.e. vertical to the arm. The arm can rotate freely about a vertical axis, by means of a free revolute joint. The natural elasticity of the arm allows for limited vertical travel. A webcam is mounted on the arm itself, next to the axis of rotation, so that it rotates in the same frame of reference as the robot, yet does not move vertically. As in the case of swinging, the webcam is used to track the position of a coloured marker placed on the robot's body. The marker's ordinate is then used as the feedback signal for the neural oscillators.

In the absence of a sufficiently large amplitude periodic feedback signal (as in the beginning of an experimental run, when the robot is stationary), the oscillator caused the motors to move at its natural frequency. This was too high for the motors that could not follow the gait trajectories in their entirety. Therefore, the steps the robot made were incomplete and stride length was small. These small and quick steps caused the robot to walk at a very slow pace. However, this caused the feedback signal amplitude to increase as the legs pushed upwards at each step, which in turn led to the oscillator entraining to the mechanical system's dynamics. This was indicated by a reduction in the frequency of the neural oscillator and significant increases in walking speed and stride length.

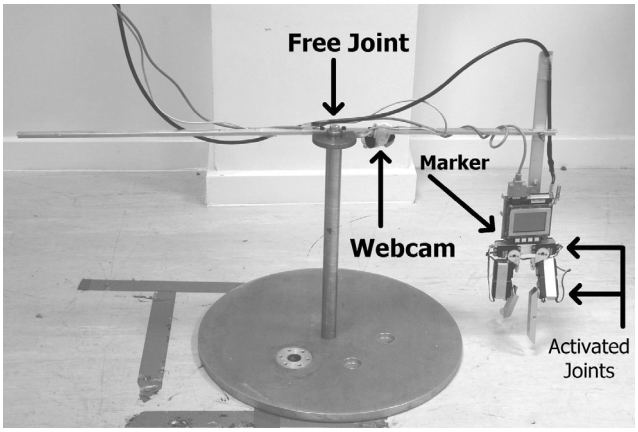

Figure 10: The experimental used for walking experiments. The elasticity of the suspending arm is exploited to absorb impact shocks and provide a rudimentary form of compliance. The vertical motion of the robot body as it walks is used as the feedback signal for the neural controller.

Data from such an experiment is shown in Figure 11, where the feedback signal, oscillator output and motor commands are plotted against (oscillator) equation time. The three arrows in the figure denote the regions where entrainment was achieved. There the oscillator produced stable, large and generally constant amplitude oscillations, while frequency- and phase-locking to the feedback signal. It becomes clear that in between these regions of stability, the oscillator's output amplitude and DC component show significant drift. However, this phenomenon is totally absent in the first case of entrainment and significantly reduced in the other two.

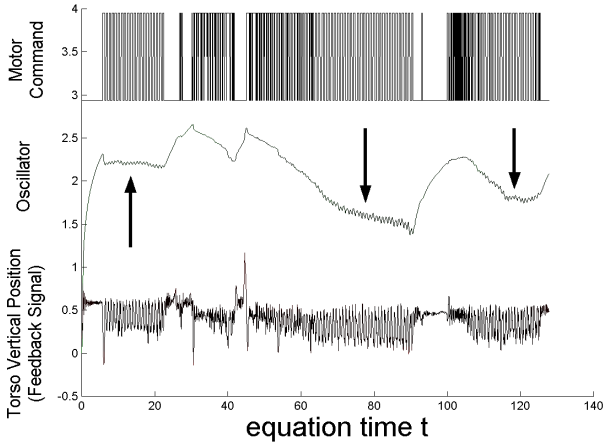

Figure 11: Time series for a typical experimental run showing, from top to bottom, the motor commands, nonlinear oscillator output and feedback signal. The arrows denote the regions where entrainment was achieved.

The phase plots for the above two signals in this experiment are shown in Figure 12. The initial high frequency oscillations occur without forming a closed limit cycle. This is possible because the motor commands are generated by taking the oscillator's derivative, rather than its direct output into account (equation 3). However, entrainment creates a stable limit cycle, seen as a 'tunnel' in the 3D plot. This type of graph additionally emphasises oscillator DC drift, allowing for a better assessment of entrainment 'quality'. Thus, in its second occurrence, the 'entrainment tunnel' can be seen drifting, indicating an imperfect lock.

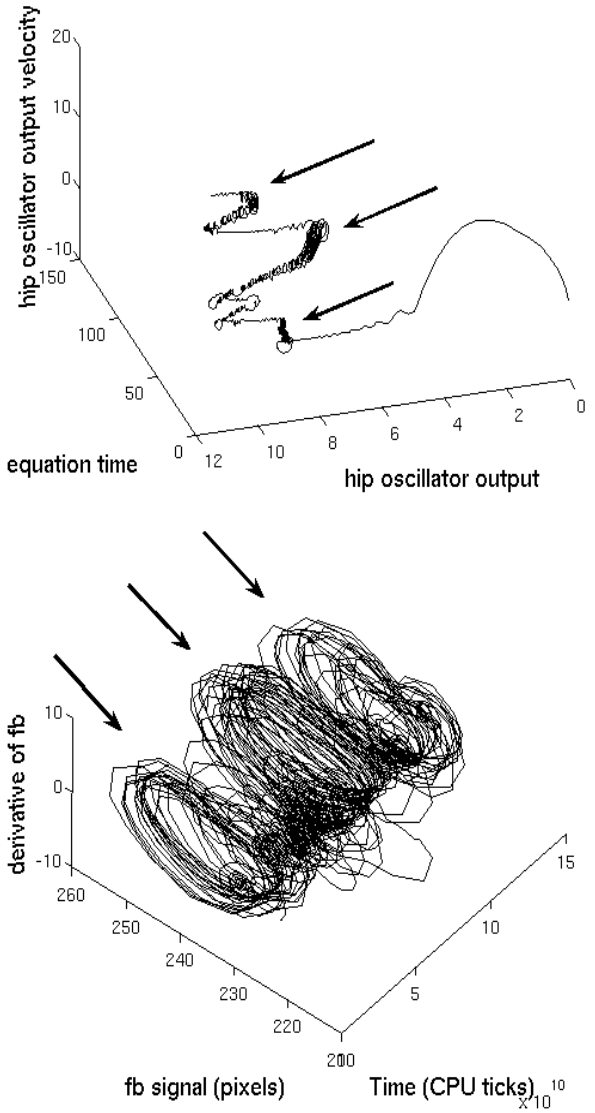

Figure 12: Oscillator output (top) and mechanical feedback signal (bottom) phase plots for the given experimental run. Again, the arrows denote the three regions during where entrainment was achieved; they are characterised by consistency in the oscillator output and large amplitude mechanical oscillations. The mechanical feedback signal has additionally been smoothed in this plot with a 5-point moving average to reduce high frequency noise and discriminate it from stable, large amplitude oscillations.

This is reflected on the mechanical signal; when the oscillator entrains well with the natural dynamics, the mechanical system increases the amplitude and velocity of its oscillations very rapidly. Since there are very few iterations with increasing amplitude, this is indicated by a clear 'entrance' to the tunnel, as in the first and third cases. In the second case, the tunnel is preceded by a growing spiral and does not maintain constant volume throughout its length. Loss of entrainment is indicated by sharply diminishing amplitude of the mechanical oscillation in all cases.

The system however, cannot maintain the entrained configuration for an extended period of time. The problem appears to be the distortion of the feedback signal by high-frequency harmonics, destroying entrainment. These harmonics are introduced after a small number of periods following the onset of entrainment. In behavioural terms, the humanoid typically makes 2-3 successful steps before motion becoming affected by these strong oscillations throughout the apparatus. They could be caused by the impact shock not being sufficiently absorbed by the compliant mechanism and thus transmitted to the robot's body.

10.Conclusions

In the robot swingin task, the van der Pol oscillator successfully managed to achieve mechanical entrainment with either the proximal or distal degrees of freedom activated. While the latter exhibits suboptimal and volatile performance on its own, enabling both DoF results in a significant improvement of performance, without affecting stability.

In the presence of strong neural entrainment, the system reaches its maximal performance. This is in agreement with prior art (Taga 1991; Lungarella and Berthouze 2002). Furthermore, the strength of the connection between the oscillators seems to be directly related to the overall stability of the system, eliminating glitches in the oscillator outputs and reducing the duration of the transient phase.

Use of the van der Pol oscillator in this study has resulted in a reduced search space compared to that of (Lungarella and Berthouze 2002), making our focus the resultant task performance characteristics instead of exploration efficiency. Our experiments demonstrate that a delayed introduction of the distal DoF does not offer any advantages over using both hip and knee from the beginning of the experiment. On the contrary, it is less efficient, both with respect to the time needed to reach a stable oscillatory regime and the maximum amplitude it can achieve.

Work is currently underway to incorporate a learning algorithm in the developmental process such as reinforcement learning (Sutton and Barto 1998). This will then allow for a more detailed exploration of the behaviour of the system and also provide insight into the interaction of the complex behaviours observed with a higher-level process.

Finally, it is of interest that a neural architecture originally designed for a robot swinging task was successful also in a robot walking task, demonstrating the elegance of the coupled dynamical systems approach.

References

Asada, M., K. F. MacDorman, et al. (2001). "Cognitive developmental robotics as a new paradigm for the design of humanoid robots." Robotics and Autonomous Systems 37(2-3): 185-193.

Beer, R. D., H. J. Chiel, et al. (1998). "Biorobotic approaches to the study of motor systems." Current Opinion in Neurobiology 8(6): 777-782.

Bernstein, N. (1967). The coordination and regulation of movements. Oxford: England, Pergamon Press.

Brooks, R. A. (1991). "Intelligence without representation." Artificial Intelligence 47: 139-160.

Goldfield, E. C. (1995). Emergent Forms: The Origins and Early Development of Human Action Systems, Oxford University Press.

Lewis, M. A., R. Etienne-Cummings, et al. (2003). "An in silico central pattern generator: silicon oscillator, coupling, entrainment, and physical computation." Biological Cybernetics 88(2): 137-151.

Lungarella, M. and L. Berthouze (2002). "On the Interplay Between Morphological, Neural, and Environmental Dynamics: A Robotic Case Study." Adaptive Behavior 10(3/4): 223-242.

Lungarella, M., G. Metta, et al. (2003). "Developmental robotics: a survey." Connection Science 15(4): 151-190.

Matsuoka, K. (1985). "Sustained oscillations generated by mutually inhibiting neurons with adaptation." Biological Cybernetics 52: 367-376.

Newell, K. M. and D. E. Vaillancourt (2001). "Dimensional change in motor learning." Human Movement Science 20(4-5): 695-715.

Strogatz, S. (2001). Nonlinear Dynamics and Chaos: With Applications to Physics, Biology, Chemistry and Engineering, Perseus Books Group.

Sutton, R. S. and A. G. Barto (1998). Reinforcement learning : an introduction. Cambridge, Mass. ; London, MIT Press.

Taga, G. (1991). "Self-organised control of bipedal locomotion by neural oscillators in unpredictable environment." Biological Cybernetics 65: 147-159.

Thelen, E. and L. Smith (1994). A Dynamic Systems Approach To The Development Of Cognition And Action. Cambridge, MA, MIT Press.

Williamson, M. (1999), Robot Arm Control Exploiting Natural Dynamics, PhD Thesis, Department of Electrical Engineering and Computer Science, Massachusetts Institute of Technology, Cambridge, Massachusetts.

Williamson, M. M. (1998). "Neural control of rhythmic arm movements." Neural Networks 11(7/8): 1379-1394.

Zielinska, T. (1996). "Coupled oscillators utilised as gait rhythm generators of a two-legged walking machine." Biological Cybernetics 74(3): 263-273.