

Assembly, Tuning, and Transfer of Action Systems in Infants and Robots

Luc Berthouze^{a,*} and Eugene C. Goldfield^b

^aCentre for Computational Neuroscience and Robotics, University of Sussex, Brighton, UK

^bChildren's Hospital Boston and Harvard Medical School, Boston, MA, USA

This paper seeks to foster a discussion on whether experiments with robots can inform theory in infant motor development and specifically (1) how the interactions among the parts of a system, including the nervous and musculoskeletal systems and the forces acting on the body, induce organizational changes in the whole, and (2) how exploratory behaviour and selective informational signals at the timescale of skill learning may allow behaviour to become stabilized at the longer timescale of development. The paper describes how three generative principles, inspired from developmental biology and shown to underlie the dynamics of infants learning to bounce in a Jolly Jumper, were broken into a set of mechanisms suitable for controlling a robotic system and resulted in a similar developmental profile. A comparison of infant and robot data leads to a set of criteria for improving the usefulness of robotic studies. Copyright © 2008 John Wiley & Sons, Ltd.

Key words: motor development; evo devo; dynamical systems; synthetic approach; infants; robots

INTRODUCTION

Papers in this special issue concern the modelling of human infant behavioural development with robots. The logic of this synthetic approach originated in part from the work of the great mathematician and Enigma code-breaker Alan Turing, in the 1950s. It is, perhaps, not well known that Turing (1952), himself, sought insights from biology and wrote an influential article on the chemical basis of morphogenesis. Turing's paper and his other work during this period stimulated the new field of mathematical biology. Turing was concerned with the mechanism by which the genes of a zygote may determine the anatomical

*Correspondence to: Luc Berthouze, Department of Informatics, University of Sussex, Brighton BN1 9QH, UK. E-mail: L.Berthouze@sussex.ac.uk

structure of the resulting organism (Keller, 2002). He proposed an influential model that has guided efforts to mathematically simulate the variety of natural forms described by biologists (see, for example, Kauffman, 1993; Winfree, 1980). Another mid-twentieth century giant in mathematics, John von Neumann, believed that the dynamics emerging from the massive iteration of simple sets of rules might generate a formal resemblance to the biological processes of self-reproduction and evolution. With the help of rapid growth in computational power of machines, the von Neumann tradition has spawned the fields of cellular automata and artificial life (a-life), and see, for example, the work of Wolfram (2002).

Now, in the early twenty-first century, there have been stunning new advances in developmental biology that have inspired new ideas in developmental psychology, especially from a dynamical systems perspective (see, e.g. Goldfield & Wolff, 2004; Thelen, Schöner, Scheier, & Smith, 2001). This paper introduces some of these new ideas and a complementary approach from the field of epigenetic robotics that together may come closer to answering some of the questions raised by Turing more than 50 years ago. In particular, we examine whether processes that govern the formation of cells into functional organs at the timescales of evolution and development may be reflected in the organization of behaviour at the timescale of microgenesis, or skill formation. Our particular focus is on the dynamics of infant motor development, and the possibility that the theory of dynamical systems and its materialization in a robot, together, may provide a methodology for answering fundamental questions such as: (1) how the interactions among the parts of a system, including the nervous and musculoskeletal systems and the forces acting on the body, induce organizational changes in the whole, and (2) how exploratory behaviour provides a means for the stabilization of some motion trajectories and not others.

Research from the newly emerging field of evolutionary developmental biology, or 'evo devo' (Carroll, 2005; Gilbert, 2003; Gilbert, Opitz, & Raff, 1996) has determined that despite the great differences in appearance, all complex animals share a common 'tool kit' of master regulatory genes that govern the formation and patterning of their body parts during embryogenesis. This tool kit now makes it possible to understand how vertebrate architecture is constructed of modular parts that differ mainly in the number and kind of repeated structures. The tool kit consists of 'on-off' switches whose exquisite timing in particular body locations is the key to answering Turing's questions about morphogenesis. The same sets of tool kit genes control the formation of body parts with similar function (but, perhaps, very different designs) among animals as different in size and form as insects and elephants (Carroll, 2005). Since these regulatory genes are shared by such disparate species, how have these differences among animals evolved? The answer is that the developmental steps executed by individual switches are interconnected into local circuits and larger networks within a genetic regulatory network architecture. These switches have the organizational property of being weakly coupled to each other and being 'poised' to rapidly switch between their stable states so that they may be easily broken or redirected for other purposes (Carroll, 2005; Kirschner & Gerhart, 2005).

Biologists have puzzled for over a century on the nature of variation, i.e. how different phenotypic properties, such as cell shape or behavioural variation, can be generated by developmental processes. An emerging view in the developmental biology is that within the cell, the generation of particular shapes depends upon internal exploratory processes guided by chemical and mechanical

information (Ingber, 2006; Kirschner & Gerhart, 2005). The exploratory elements within the cell take the form of microtubules that constantly grow and shrink, and as they form particular arrangements under the guidance of chemical gradients and mechanical forces, a particular cell shape takes form. Here, we explore the possibility that at the macroscopic scale of behaviour by animals, including humans, similar exploratory processes may govern the assembly of stable functional units.

We propose a set of principles from developmental biology for understanding the processes of exploration and selection in assembling new behavioural forms as follows:

- (1) *Assembly from existing parts.* The development of new forms is made possible by the re-use (dissolution and re-organization) of a set of component parts.
- (2) *Exploration.* Developmental variation is made possible by exploratory processes guided by informational signals capable of stabilizing the component parts into temporary functional units.
- (3) *Selective stabilization.* The functional units of a system are assembled and dissolved under the influence of information-based selective processes that stabilize some patterns and not others.

INFANT MOTOR DEVELOPMENT: ASSEMBLY AND SELECTIVE STABILIZATION BY MEANS OF EXPLORATION

We next turn to the question whether assembly from existing parts, exploration, and selective stabilization, as shown by scientists in the field of developmental biology, may also be apparent in the organization of motor behaviour as infants explore their own activity. Is there evidence that functional units of motor behaviour are assembled and dissolved as a consequence of active exploration? The strategic use of humanoid robots to embody the underlying dynamics proposed for motor development may help structure the evidence from studies of infant behaviour into a consistent theoretical framework. To begin to address this possibility, we (a) present an outline of a dynamical systems perspective on infant motor development, (b) summarize research conducted with human infants on motor development that supports the underlying principles of exploration and assembly by selective stabilization, and (c) examine robotic data from a synthetic approach in support of the experimental evidence from human infants.

Goldfield (1995) and more recently Goldfield and Wolff (2004) and Goldfield (2007) proposed that the functional units of infant motor behaviour are assembled from components of the body's neuromuscular circuits and networks that may span several organ systems. These functional units may be weakly coupled so that they may be rapidly assembled, dissolved, and then re-assembled in novel ways under the selective stabilization of receptor information. So, for example, under conditions when the goal is to bounce up and down, the legs may be made to act 'springy', and the body 'becomes' a bouncer. Placed on a treadmill, by contrast, the legs become organized into a pendular system and become 'steppers'.

Here, we extend the theory by proposing a dynamical toolkit for assembling and dissolving temporarily assembled organizational units, abstractly the parameters and variables of a mathematical function generating the observable

dynamics. Below, we present infant data to suggest, and robotic data to further investigate, the nature of the exploratory process by which the more macroscopic functional units are assembled, rapidly dissolved, and reformed. By systematically varying putative control parameters of the temporarily assembled system, we show that it is possible to drive the system through different stable states (e.g. two configurations of a bistable system).

Goldfield, Kay, and Warren (1993) and Foo, Goldfield, Kay, and Warren (2004) carried out a longitudinal study of eight 6-month-old infants placed in a 'Jolly Jumper' infant bouncer (a harness suspended from a linear spring), with the soles of the feet just touching the floor. There were no instructions or models, and it was left to the infants to discover the consequences of kicking against the floor. The authors addressed a particular question about the assembly of a motor skill during infancy, namely how infants discovered that their spontaneous kicking against the ground eventually resulted in sustained 'bouncing'. Placed in the context of the aforementioned toolkit, the question takes the following form: how does spontaneous activity induce the macroscopic degrees of freedom (DOFs) of the legs to become a periodic 'kicker' with limit cycle, or sustained oscillatory, behaviour.

The discussion above, and the remainder of this paper, uses the word 'exploration' in ways that may not be consistent with the typically accepted definition of the word. It may therefore be useful to clarify what is meant here.

Exploration Versus Goal-directed Exploratory Behaviour

A distinction may be made between exploration and goal-directed exploratory behaviour. The distinction follows from the dynamical view that the behaviour can be viewed at (a) a lower level of local interactions between agent and environment, and (b) global structure at an abstract level of attractor dynamics (see, e.g. Mayer-Kress, Liu, & Newell, 2006; Warren, 2006). Exploratory behaviour viewed at the local level of perceiving and acting consists of an agent performing an action and exerting forces with the body effectors that alter the state of the environment. A consequence is that new information is generated about the current state of the agent–environment system, and perceiving and acting progress through another cycle. Hence, for example, when an infant manipulates a toy and observes changes in its surface relationships due to rotation and translation, there is a local cycle that promotes opportunities for further exploration.

Exploration viewed at the global level of adaptive behaviour is governed by the time evolution of stable states, called attractors and repellers. Each goal at this global level may be described abstractly as a region in a mathematical state space towards which behavioural trajectories converge (attractors) or diverge (repellers), with transitions between regions leading to abrupt changes in the trajectory. The stability of each goal state evolves and dissolves according to a particular time evolution, i.e. at particular scale of time. Behavioural development of an individual human is envisioned as being governed by a landscape of attractors and repellers that evolve and dissolve over time. As an attractor evolves and dissolves, its stability changes concomitantly. A critical point is that attractors and repellers at the slower timescale of ontogeny dominate the faster timescale of perception–action cycles, and exert stabilizing and destabilizing influences on it. Goal-directed exploration occurs when the behaviour may be shown to be dominated by the slower timescale of ontogeny and, conversely, exploration in

the service of local cycles of perceiving and acting may be apparent only at the faster timescale of learning.

One way to illustrate this distinction is to examine the influences of ontogenetic predispositions in kicking coordination on the learning of a new way to kick in order to make an interesting sight occur. In a series of studies, early work by Thelen (1985) showed that 3-month-olds preferred to kick either with both legs in alternation or with a single leg. Thelen (1994) used this finding to determine how the ontogenetically determined preferred pattern of kicking would be modified by learning to kick in a different way during a conjugate reinforcement paradigm. In this task, infants learned to kick in order to make a mobile move. The infant's legs were loosely yoked together with a tether so that if one leg kicked, it pulled on the other leg. Thus, the most efficient way to get the mobile to move was to simultaneously kick both legs (again, not the preferred kicking pattern at this age). In the acquisition phase of the task, Thelen found that infants rapidly learned to simultaneously kick the legs to make the mobile move. However, during the extinction phase when the yoke between the legs was removed, the infants quickly lost the simultaneous coordination pattern. In other words, as soon as the local conditions of the task for kicking with both legs was removed, the more global tendency for kicking with one leg again emerged. In the discussion below of another task involving infant kicking, bouncing, we would argue that exploration during a particular learning session reflects the local interactions between body and environment during those fleeting moments when infants are kicking and experiencing a consequential body motion. By contrast, we believe that the exploratory behaviour in the context of changing behaviour over longitudinal sessions is more likely to be indicative of attraction towards an ontogenetic goal.

On this view, the model presented below has two particular requirements owing to the characteristics of animals in their environment. First, the functions used to model the behaviour should exhibit goal directedness, i.e. the tendency to arrive at some anticipated endpoint, or goal. Second, the functions should capture the biomechanical properties of the body as they react to contact with the substances and surfaces of the environment. We will return in the final section to the question whether synthetic analyses using humanoid robots to simulate infant behaviour satisfy these two requirements.

A Toolkit for Microgenesis

Goldfield *et al.* and Foo *et al.* began their analysis by modelling the goal of the system, stable bouncing, with a function that exhibits a particular attractor dynamics, namely a tunable forced mass–spring system:

$$m\ddot{x} + b\dot{x} + kx = F_0 \cos(\omega t)$$

where x is the position of the mass m , a single dot above x indicates velocity, and a double dot indicates acceleration, k is the spring stiffness, b is a damping coefficient or friction term, and F_0 is an external force that occurs with frequency ω at an external timescale t . A mass–spring system that is 'forced' is one that includes some regular periodic source of energy input. The parameters of the system, mass (m), stiffness (k), and damping (b) when assembled in this way represent the properties of the physical spring on the left side of the equation, and the infant's kicking on the right side. However, while the right side of the equation indicates continuous sinusoidal functioning, during sustained

bouncing, the infant's feet are actually in contact with the ground for less than half the cycle. During contact with the ground, the legs act more like springs than pure force applicators, with the joints and muscles having stiffness and damping characteristics of their own.

Further consideration of the biological properties of the legs suggests that the equation must include not only the relation of kicking to the timing of contact with the ground, but also the transmission of force between the muscles and the mass of the infant's body. Maximum power can be transmitted to the load if the impedance properties of the transmission (the muscles) are matched to the impedance properties of the load (the body). Another consideration of placing the body in the context of the support surface is the information available to the infant for the appropriate frequency and phasing of kicking (i.e. when to kick in order to keep bouncing) that may come from the foot contact with the ground. Haptic information from the foot during the contact phase of bouncing transforms the above equation from a linear externally driven mass-spring into an autonomous limit cycle (i.e. oscillatory) system with the intrinsic timing modulated by foot contact with the surface, a property of non-linear oscillators. Therefore, Goldfield *et al.* proposed the following function as a model of infant bouncing:

$$m\ddot{x} + b\dot{x} + (k_{\text{spring}} + k_{\text{legs}})x = F(\Phi)$$

Now, the right side of the equation, the macroscopic characteristic of frequency and phasing of kicking, is some function of the phase of the body mass's motion, and the left side of the equation includes stiffness contributions of both the physical spring and the infant's legs.

A fundamental characteristic of the parameters comprising the assembled function is that they are tunable. So, for example, stiffness may assume a range of values as the system explores its own behaviour. By making the system parameters tunable, it becomes possible to capture the process by which changing certain inputs to the assembled macroscopic system results in a modification of bouncing behaviour, without decomposing the organization of the behaviour, itself. That is, the function representing the behaviour may yield different trajectories (behavioural forms) as a consequence of a modification of its parameters (experience), but without losing its integrity, i.e. it still describes the behaviour and its changes. For example, once infants begin to bounce, they may explore their own behaviour by varying one of the parameters while holding the others constant. The emphasis here is on the exploratory nature of tuning parameters, making tuning a process whose goal is to find in the space of all possible trajectories the attractor state that defines an operation at the resonant frequency (see below). The system may tend towards a goal of kicking at the precise moment in the cycle of the spring that captures its energy and sustains a stable oscillation. In physical terms, for a given kicking force, the amplitude of resulting oscillations is maximal at a specific frequency, its resonant frequency. Since resonant frequency depends upon stiffness of the spring and legs, infants may learn about the relation between their muscular efforts (e.g. when to kick) and the consequences for body motion as system stiffness is changed. As a goal-directed process, exploration may serve the function of selectively stabilizing certain bouncing trajectories and eliminating others, under the influence of the haptic input deriving from the stimulation at the sole of the feet. If so, then as infant gains experience in bouncing, initial variability in kicking frequency and stiffness should gradually decrease.

Hypothesis Testing with Babies

The following hypotheses and the accompanying evaluation, tested in longitudinal studies by Goldfield *et al.* and by Foo *et al.*, address the process of assembly from the existing parts and selective stabilization (tuning) by means of exploratory behaviour:

1. There should be an early assembly phase characterized by sporadic, irregular kicking without sustained bouncing. In both studies, during early sessions infants kicked irregularly with only one or two successive bounces before a pause (called a 'bout'). Bout length increased over several sessions until it suddenly doubled in a peak session.
2. Emerging from this should be a 'tuning' phase with more periodic kicking, during which the frequency of flexion–extension cycles and leg stiffness vary, yielding initial high variability in the bouncing period. Goldfield *et al.* and Foo *et al.* found a steady decline in the variability of period, resulting perhaps from the increased use of foot contact information, which intrinsically specifies frequency, and adjusting leg stiffness to match spring stiffness.
3. Once bouncing stabilizes at its limit-cycle attractor, i.e. occurs with regular repetitive cycles, a sustained bouncing phase should occur with the following characteristics: (a) oscillation at the resonant period, (b) decrease in the variability of period, and (c) increase in the amplitude due to operating at resonance. Goldfield *et al.* and Foo *et al.* found that the infant's preferred period of oscillation closely approximated the resonant period of the system. Moreover, there was a decreased variability in the period during sustained bouncing, as would be expected if the system reached its goal of achieving stability. Such stability is, again, characteristic of a system using haptic information to regulate forcing frequency. And, as predicted for a system at resonance, the amplitude of bouncing increased dramatically at the 'peak' session.
4. If the infant has learned the low-dimensional dynamics of the task (i.e. the macroscopic behaviour of the legs as a 'bouncer') rather than a specific forcing frequency and leg stiffness, there should be a rapid adaptation to changes in the task conditions. When Foo *et al.* experimentally modified the task conditions by introducing a spring with different stiffness after the infant had experience with an initial spring, there was a rapid return to the same bouncing behaviour.

These four major findings suggest that infants actively explore their own behaviour via a process of variation, and that initially variable and disorganized actions become stabilized into useful forms as a consequence of tuning the parameters of the toolkit components.

ROBOTIC MODELLING OF INFANT BOUNCING

The premise of the synthetic approach is that designing an artificial system to perform a particular behaviour should be based upon an attempt to use the same underlying principles to build robot architectures as those hypothesized to govern infant behaviour. So, for example, the toolkit of mathematical functions with tunable parameters should be able to generate the same behaviour, whether instantiated in infant tissue or artificial materials. It is not simply a matter of

showing that a robot can perform the same observable behaviours. Rather, it is by virtue of using the same underlying generative principles, the toolkit of mathematical functions from work in dynamical systems, that any similarities between robot and baby have meaning.

The synthetic process of 'understanding by building' has met with some success, particularly in invertebrate neuroscience where robots have been used to test hypotheses and make predictions about various aspects of sensorimotor control (see Webb, 2002 for a review). However, there has been only limited work using robots to examine the complex interplay between body, neural systems, and environment as babies explore ways to use their bodies to achieve particular goals.

Given these caveats about the significance of the similarity between generative systems in baby and robot, this section briefly introduces a research program that uses robots to further examine questions raised by studies of infant bouncing. These include:

- (1) What is the nature of the functional units, and at what level of abstraction should they be implemented?
- (2) What are the environmental and biomechanical variables of interest? Indeed, the choice of those variables has implications for the design of both the sensory and mechanical systems of the robot.
- (3) How much prior knowledge (task, environment, body structure) is available to the child/system learning to bounce? And what should be a suitable model of the goal and motivation of the child/system?

Assembly of Organized Behaviour

In a series of experiments, Berthouze and colleagues set out to reproduce the longitudinal profile of learning to bounce by using the synthetic approach (Lungarella & Berthouze, 2004; Spröwitz & Berthouze, 2005). They constructed a small biped robot with three actuated DOFs per leg. The robot was strapped in a Jolly Jumper, and its movement constrained to the sagittal plane by a parallel guide. Four force sensing resistor sensors are placed underneath its feet and an accelerometer provided the sensory feedback. The feet as well as the joints were designed to be passively compliant (Meyer, Spröwitz, & Berthouze, 2006). The elementary functional units were modelled as Bonhoeffer—van der Pol neural oscillators (Fitzhugh, 1961). This formalism is a simplification of the variable Hodgkin—Huxley model to a simpler algebraic form with two variables (excitable and recovery variables) and intended to represent the qualitative properties of a wide class of such oscillators. With its ability to produce oscillatory behaviour, the neural oscillator satisfies the requirement discussed earlier that the function used to model the behaviour should exhibit goal directedness, or the tendency to arrive at some anticipated endpoint, or goal. Unlike harmonic oscillators, Bonhoeffer—van der Pol oscillators exhibit robust phase locking, or entrainment, over a wide range of time delays, which makes them particularly attractive to study ideas such as assembly, tuning, and sensitivity to environmental input. Finally, assuming a suitable connectivity matrix, a collection of oscillators provides enough flexibility to study different tasks. And indeed, central pattern generators (CPGs) are found to subservise a variety of motor acts in vertebrates, including humans (e.g. Grillner, 1985; Grillner, Markram, De Schutter, Silberberg,

& Le Beau, 2005). Furthermore, whilst CPGs are traditionally studied in the context of rhythmical activities, a recent study by Schaal, Sternad, Osu, and Kawato (2004) also showed that brain areas involved in the generation of rhythmical activity subserve the generation of discrete movements.

In this oscillator-based framework, Berthouze and colleagues suggest that the assembly phase (as defined in Section 'Hypothesis testing with babies') can be interpreted as the self-organization of the collection of oscillators—the adaptation of the connection weights between each oscillator—and the tuning phase as the time-constant adjustment for each oscillator. The synthetic experiments examined the role of exploration by determining whether a robot capable of modifying its behaviour on the basis of sensor data would begin to organize its behaviour in a way similar to the longitudinal pattern observed in human infants. A goal of these experiments was to determine whether exploratory behaviour could reasonably be characterized as the search for optimal task-specific sensorimotor configurations under a condition of information maximization. The results of such experiments may provide a new impetus to building robotic architectures that take best advantage of autonomous exploratory behaviour (see Fitzpatrick, Needham, Natale, & Metta, 2008, this issue).

Searching the Parameter Space

Lungarella and Berthouze (2002), for example, modelled the exploratory activity of a robotic system as a directed random search of the parameter space, using a Metropolis-like search (Metropolis, Rosenbluth, Rosenbluth, Teller, & Teller, 1953), biased by a value system based on two qualities of the resulting bouncing behaviour, namely stability and height. The trade-off between exploration (of the parameter space) and exploitation (of each configuration) was implemented by controlling the rate of randomness of the search according to an annealing-like process. Simply stated, simulated annealing (Kirkpatrick, Gelatt, & Vecchi, 1983) is a stochastic optimization method for finding global minima in a large search space, inspired by the metallurgic process whereby a material is repeatedly heated and then slowly cooled to enable its atoms to leave a particular state, wander randomly through states of higher energy, and converge onto configurations with lower energy than the initial one, i.e. more stable states. In the context of this study, exploration (change in the parameter space) is facilitated when behavioural performance is low (high temperature) and conversely, exploitation is favoured when a particular parameter configuration yields higher performance, in which case the changes in parameters become smaller (low temperature). In the value space, i.e. the space that maps each set of parameters to its resulting behavioural score or value, this process results in a system starting with a pseudo-random trajectory and eventually reaching an attractor state from which the system will only leave if it is perturbed or if habituation (here modelled as a reset of the temperature, re-annealing) kicks in (see Figure 1).

In Berthouze *et al.*'s experiments, the robot's behaviour was characterized by an initial phase without bouncing, i.e. without distinct vertical flight and descent, and followed by a sporadic low-amplitude bouncing at which point the tuning phase was triggered. An important characteristic of this latter phase was the occurrence of multiple stable formations (different joint synergies yielding similar trajectories, see Figure 3 for some examples of ankle–hip phase plots), the importance of which will be discussed in the following section. The profile of the bouncing height during the tuning phase showed a bell-shaped curve

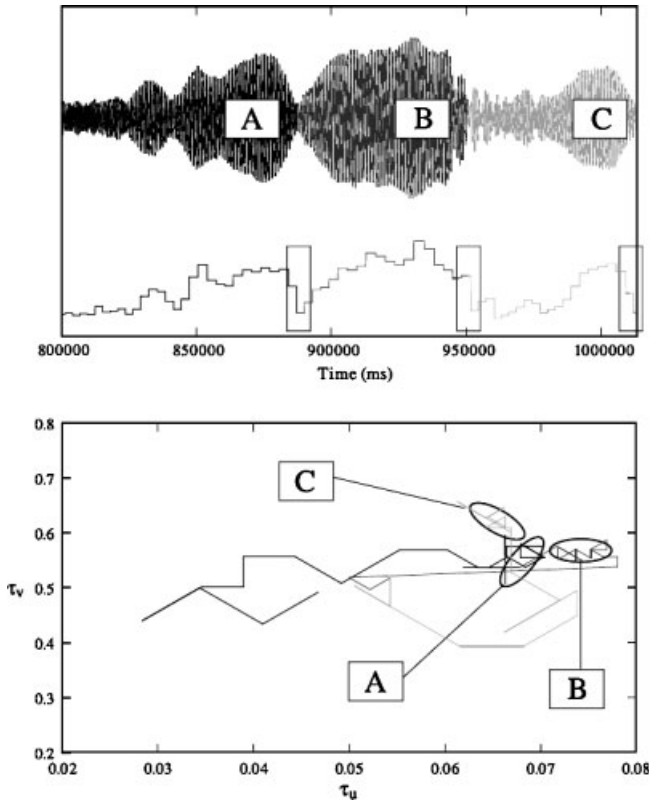


Figure 1. The upper graph depicts the time series of the oscillatory movement of the robot's hip (top) and the associated behavioural score, the inverse of which controlled the temperature parameter. Rectangular areas point to decreases of the score by habituation (re-annealing). The lower graph depicts the corresponding trajectories in the parameter space (here, two parameters controlling the timing of a particular oscillator). Oval areas denote dense regions of high-yield parameter settings, that is, large oscillations observed in the time series (figure taken from Lungarella, & Berthouze, 2002).

qualitatively similar to that predicted by Goldfield, with quantitative differences reasonably well explained by the physical characteristics of the springs (spring constant and damping coefficient) and the system (mass, compliance). Among those characteristics, compliance was found to play a very important role by not only allowing for lower torques, but also, and most importantly, by improving the quality of the sensory signals (low-pass filtering) which, in turn, facilitated entrainment.

Further experiments have examined whether the emerging configuration of robotic motor behaviour was task specific, i.e. rapidly adaptive to a change in a task variable, such as load. For example, Spröwitz and Berthouze (2005) simulated a morphological change in the robot, by keeping the control parameters constant and by changing the robot's lift-off weight. Despite the change in dynamic load, the robot showed seamless adaptation to the changes, by way of entrainment made possible by a stable configuration of the functional units (see Figure 2). This is similar to what was reported by Foo *et al.* with infants when they changed the spring constant.




Figure 2. Time series of interbounce intervals when the robot's (lift-off) weight is varied (cyan, orange, and green areas). Purple areas denote periods of manual perturbations. All axes are in milliseconds. This figure is available in colour online at www.interscience.wiley.com/journal/icd.

Optimal Joint Synergies: Hard-wired or Self-organized?

The occurrence of multiple stable formations within the assembly phase leads to the question of whether joint synergies should be seen as a result of hard-wired connectivity or, instead, as a dynamically emerging phenomenon. In two studies, Lungarella and Berthouze (2002) and Berthouze and Lungarella (2004) revisited the notion of free(z)ing of DOFs as a solution to the so-called 'DOF problem' (Bernstein, 1967), namely the fact that more than one motor signal can lead to the same trajectory of a given motor system (the redundancy of effective movements). Bernstein argues that in the early stages of learning a new motor skill, the complexity of the task is initially reduced by freezing out some of the DOFs (often, the distal ones), either by keeping them rigidly fixed, or by tightly coupling them (see Vereijken, Whiting, Van Emmerik, & Newell, 1992, for example). The frozen DOFs are then subsequently released.

Recent developments in ecological psychology and dynamical systems suggest that freezing and freeing might pertain to the dynamics of the developing motor system (e.g. Sporns & Edelman, 1993). Here, we question the role of this phenomenon, especially in terms of how it might act as a stabilizing mechanism (principle 3 in the Introduction) by structuring the way information is gained during exploration or learning. Elman (1993), for example, shows how artificial neural networks can only learn to process complex grammars when they are handicapped by severe capacity limitations, i.e. by reducing the available DOFs. Berthouze and Kuniyoshi (1998) use freezing as a means of simplifying the acquisition of complex visuomotor coordination in robots with multiple redundant DOFs.

Through the analysis of another rhythmical stereotypy, swinging, Berthouze and colleagues ask two questions: (a) what, if any, is the gain of freezing? And (b) is that gain/benefit robust to unexpected perturbations and increase in task complexity? Swinging is an interesting case study because it is difficult to solve from a control point of view (double inverted pendulum) and optimal task performance requires efficient joint synergies and accurate (anticipatory) timings. Using their biped robot attached to a freely rotating joint, Lungarella and Berthouze (2002) proposed a comparative analysis of a systematic search of the space of parameters under three different conditions: (a) when all DOFs are used, (b) when the distal DOFs are frozen, and (c) when the distal DOFs are released after a stationary regime is obtained in the frozen configuration.

The outright use of all DOFs was characterized by a rugged landscape with low value. Task performance was poor (low amplitude, little or no coordination,

etc.) and qualitatively comparable to what was being observed during the assembly phase of the infant bouncing study. Parameters taken in a small neighbourhood yielded qualitatively different behaviours and multiple stationary regimes co-existed within a same configuration (see Figure 3).

Freezing the distal DOFs, instead, resulted in a large basin of attraction (see Figure 4) where large amplitude oscillations were obtained independently of the initial conditions, a finding that can be attributed to physical entrainment.

In an empirical confirmation of Bernstein's idea, the release of the distal DOFs after stationary regime was attained in the frozen configuration led to another value landscape in which all initial conditions yielded smooth large amplitude oscillations, i.e. the temporary reduction in the number of DOFs provided for a more efficient exploration of the space of parameters. A similar result was obtained when the reduction of the number of DOFs was realized by way of synergies. Optimal performance was achieved when the coupling gain between oscillators was strong enough for neural entrainment to take place and to stabilize the system in motion (emergence of a single flexible DOF). In a subsequent study, Berthouze and Lungarella (2004) tested the robustness of that mechanism to an increase in task complexity, obtained by adding an asymmetric non-linear coupling between system and environment. The study revealed that a single pathway of freezing followed by freeing, as used in the previous study, was not sufficient to yield stable behaviour, but that, instead, alternate phases of freezing and freeing were necessary.

The results, thus, support the idea that free(z)ing should be seen as a dynamic process with its own (low) time constant that stabilizes faster adaptation

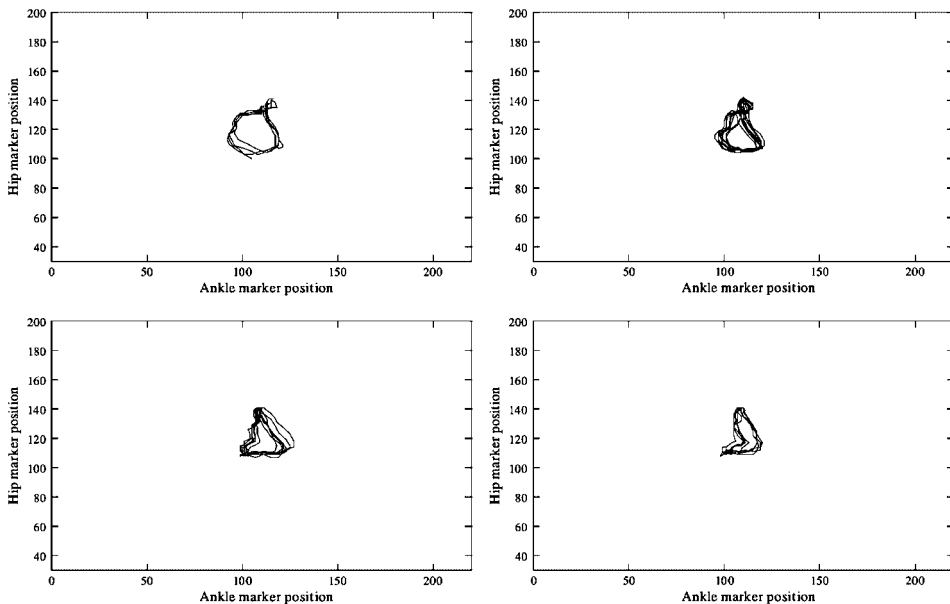


Figure 3. Existence of preferred stable states and phase transitions in the independent 2-DOF configuration. The graphs show the ankle-hip phase plots of four pseudo-stationary regimes obtained within a single parameter configuration (called multiple stable formations in the text).




Figure 4. Large basin of attraction uncovered by exploration of the parameter space with a frozen distal degree of freedom. The size of a dot (a control setting defined by the two-parameter settings in the x and y coordinates) is proportional to the behavioural score obtained for that particular control setting. This particular explorative run took roughly 10 min (figure taken from Lungarella & Berthouze, 2002).

processes (e.g. neural entrainment) and can push the system in, or away, from low-value configurations. In a recent study, Hong and Newell (2006) report findings in which 'differential changes in the organization of system properties with learning are consistent with the proposition of Berthouze and Lungarella that alternating freezing and freeing of mechanical DOFs serves as the ideal strategy of change for the achievement of skill acquisition'.

CONCLUSIONS

This concluding section returns to the question of whether synthetic experiments with robots may increase our understanding of infant motor development. The question centres on the criteria for determining how well the robot used in experiments by Berthouze and colleagues simulated infant motor development, with specific reference to bouncing. We began the paper with Turing's early contributions to mathematical inquiries into growth of biological systems. Turing was, of course, one of the first to propose a simulation methodology for evaluating whether a machine could think. His imitation game, famously known as the Turing test, consisted of a set of questions by a hidden machine, designed to fool a human subject into thinking that its responses were produced by another person. More than 50 years later, scientists in emerging fields, such as developmental robotics (Lungarella, Metta, Pfeifer, & Sandini, 2003), biorobotics (Webb, 2001), and ecological robotics (Duchon, Warren, & Kaelbling, 1998) have

proposed new tests for deciding whether particular robotic behaviour serves as a good model of biological behaviour. These new tests go well beyond Turing's disembodied tickertape. Here, we highlight three criteria for evaluating robotic behaviour: embodiment, development, and selective use of available multimodal information.

Table 1 summarizes our attempt to use specific success criteria to evaluate how well the behaviour of a robot modelled the development of infant bouncing. The first column of Table 1 again presents the three processes proposed above to explain the observed developmental changes in infant bouncing behaviour. A particular success criterion is used to evaluate the robotic modelling of assembly, tuning, and stabilization, respectively. The first process, assembly, addresses the nature of the functional units that are temporarily assembled for the particular task of bouncing. The success criterion is called embodiment, the extent to which the functional unit includes both neural oscillators and a physical body governed by the forces and fields of the environment.

The Berthouze *et al.* robot was designed explicitly to mimic the way in which neural oscillators are coupled to compliant mechanical systems. One major difference between changes in the robotic behaviour and the development of human infants is that the robot was designed for the specific task of bouncing, and is unable to use the same functional units for bouncing in order to locomote, even by crawling or dragging itself across the floor as babies do. For example, Goldfield (1989, 1993) (but see also more recent work by Adolph, Vereijken, & Denny, 1998) demonstrated that stages in the development of an infant locomotor skill, crawling, reflected the way that the particular functional capabilities of the eye-head system, hands, and legs became assembled into a multifunctional system. Goldfield proposed that particular parts of the anatomy are temporarily coupled together as functional units, depending upon the task. Crawling, for example, may consist of the legs, hands, and eye-head system temporarily acting together for the organized activity of propulsion, steering, and balance, respectively. Goldfield found that stages in the development of crawling were best characterized not by a lock-step change in flexion-extension patterns of muscle groups, but rather, were an emergent outcome of the way that different parts of the anatomy were coupled together cooperatively. During a portion of the locomotor cycle of crawling, the hands are used for support, but during

Table 1. Evaluating robot bouncing and its development

Process	Baby	Robot	Success criterion
Assembly	Putative coupling between neural oscillators and compliant body	Coupling of van der Pol oscillators to passively compliant body	Embodiment (interaction between neural oscillators and compliant body)
Tuning	Infants learn to kick at resonant period of system in order to take advantage of compliant body and spring	Explicit use of random search and optimization process	Developmental progression from sporadic to sustained bouncing
Stabilization	Sustained bouncing occurs at resonant period of the system; decreased period variability; increased bounce amplitude	Experiments demonstrate stabilization (freezing and freeing) of degrees of freedom	Selective use of available information to stabilize some trajectories and not others

another part of the cycle, they are used for reaching ahead to the next surface location used for support. By coupling the propulsion function of kicking to the balance function of the eye-head (visual-vestibular) system and the steering function of the two hands on the support surface, the infant's efforts at forward progression resulted in a mode of activity that integrated multiple functions towards the achievement of a unified goal, namely stable locomotion in a particular direction. At this point, robots have only limited capacity to use the same parts for different functions (see Platt, Fagg, & Grupen, 2003, for a rare exception), and a more comprehensive set of synthetic experiments will await further advances in robotics.

The second process, tuning, is explicitly developmental because it is defined by changes in some outcome measure over time. For example, Goldfield *et al.* (1993) and Foo *et al.* (2004) found that infants progressed from sporadic kicks to smooth and effortless bouncing. The Berthouze *et al.* robot, too, produced a change in bouncing behaviour over time. As with babies, the behaviour of the robot was shown to not be due simply to a specific forcing frequency and stiffness, but rather, reflected the low-dimensional dynamics of the task. But does this mean that development occurred in the robot's behaviour? An obvious difference from human infants is that any robot built similar to the one used by Berthouze would look the same over time. This stands in marked contrast to the development of any two babies. There was a considerable range between babies in the number of days required to progress to sustained bouncing, and these differences reflected a variety of factors including rates of body growth and muscle strength. Even though the endpoint was similar for a robot and for infants, the process of getting there was quite different.

The third process, selective stabilization by means of exploratory behaviour, incorporates environmental information into the control process. Both the infants and robot provide evidence that the emergence of long sequences of high-amplitude bouncing was made possible through discovering the resonant period of the system. However, use of the sensory information available during bouncing by the robot was quite different from that of infants. Surface contact between the robot and environment was limited to forcing resistors under its feet, and body motion to an accelerometer. By contrast, the active exploratory nature of the infant's eyes, ears, and haptic perceptual systems provides a rich, multimodal source of information about their motion with respect to the support surface and to their surroundings. This information is available to infants' nervous systems as a means for stabilizing their motion. Berthouze *et al.* have demonstrated a process of freezing and freeing of DOFs in robot swinging. This implies that the robotic control system may be characterized as a low-dimensional system with a few free control variables (Warren, 2006). However, there is not yet any specification of the information available to the robot that may map onto these control variables, such as that created during vertical flight and descent.

Our suggestion that experiments with robots could inform theory in infant motor development may, at first sight, appear premature or even presumptuous. Modelling infant motor development with robots present technological challenges (that may or may not be solved in the short term) as well as methodological challenges. Still, we believe that this work shows the promise of using embodied agents, as identified by Webb (2001) and others, that is, to go from a mere description of putative mechanisms to a physical instantiation of these mechanisms that will allow for the theory to be put to the test. This is only possible if 'structural validity' is achieved, i.e. when the model not only produces

the observed real system behaviour but also truly reflects the way in which the real system operates to produce this behaviour (Zeigler, 1976). In the work described above, (a) the same toolkit of mathematical functions was used to generate a large set of potential solutions to the problem of bouncing that characterized both infant and robot behaviour, and (b) the organismic and environmental tuning of the parameters of these functions made it possible for the infant and robot to discover ways to harness the properties of both muscles and actuators, respectively.

Identifying those tasks and mechanisms that afford structural validity will be an important key for robotic experiments to inform theory in infant motor development and will benefit most from a reinforced interaction between our disciplines.

ACKNOWLEDGEMENTS

L. Berthouze was supported by a Grant-in-Aid for Young Scientists (no. 15700174) from the Japanese Ministry of Education, Culture, Sports Science and Technology (MEXT). E. C. Goldfield is supported by the following grants from the US National Institutes of Health: 1 R01 DC007127, 1 R43 HD047128, 1 R43 HD047128, and the Developmental Disabilities Research Center at Children's Hospital Boston, P30 HD18655.

REFERENCES

- Adolph, K. E., Vereijken, B., & Denny, M. A. (1998). Learning to crawl. *Child Development*, 69, 1299–1312.
- Bernstein, N. (1967). *The coordination and regulation of movements*. Oxford: Pergamon.
- Berthouze, L., & Kuniyoshi, Y. (1998). Emergence and categorization of coordinated visual behaviour through embodied interaction. *Machine Learning*, 31(1–3), 187–200.
- Berthouze, L., & Lungarella, M. (2004). Motor skill acquisition under environmental perturbations: On the necessity of alternate freezing and freeing of degrees of freedom. *Adaptive Behavior*, 12(1), 47–63.
- Carroll, S. B. (2005). *Endless forms most beautiful: The new science of evo devo*. New York: Norton.
- Duchon, A. P., Warren, W. H., Jr., & Kaelbling, L. P. (1998). Ecological robotics. *Adaptive Behaviour*, 6, 473–507.
- Elman, J. L. (1993). Learning and development in neural networks: The importance of starting small. *Cognition*, 48, 71–99.
- Fitzhugh, R. (1961). Impulses and physiological states in theoretical models of nerve membrane. *Biophysical Journal*, 1, 445–466.
- Fitzpatrick, P., Needham, A., Natale, L., & Metta, G. (2008). Shared challenges in object perception for robots and infants. *Infant and Child Development*, 17, 7–24.
- Foo, P., Goldfield, E. C., Kay, B., & Warren, W. H., Jr. (2004). Infant bouncing: Spontaneous learning of task dynamics. *Proceedings of the fourteenth international conference on infant studies*, Chicago.
- Gilbert, S. F. (2003). *Developmental biology* (7th ed.). Sunderland, MA: Sinauer.
- Gilbert, S. F., Opitz, J. M., & Raff, R. A. (1996). Resynthesizing evolutionary and developmental biology. *Developmental Biology*, 173, 357–372.
- Goldfield, E. C. (1989). Transition from rocking to crawling: Postural constraints on infant movement. *Developmental Psychology*, 25, 913–919.
- Goldfield, E. C. (1993). Dynamic systems in development: Action systems. In E. Thelen & L. Smith (Eds), *Dynamic systems in development: Applications* (pp. 51–70). Cambridge, MA: MIT Press.

- Goldfield, E. C. (1995). *Emergent forms: Origins and early development of human action and perception*. New York: Oxford University Press.
- Goldfield, E. C. (2007). A dynamical systems approach to infant oral feeding and dysphagia: From model system to therapeutic medical device. *Ecological Psychology*, 19(1), 21–48.
- Goldfield, E. C., Kay, B., & Warren, W. H., Jr. (1993). Infant bouncing: The assembly and tuning of an action system. *Child Development*, 64, 1128–1142.
- Goldfield, E. C., & Wolff, P. H. (2004). A dynamical systems perspective on infant action and its development. In G. Bremner & A. Slater (Eds.), *Theories of infant development*. Oxford: Blackwell.
- Grillner, S. (1985). Neurobiological bases of rhythmic motor acts in vertebrates. *Science*, 228, 143–149.
- Grillner, S., Markram, H., De Schutter, E., Silberberg, G., & Le Beau, F. (2005). Microcircuits in action—from CPGs to neocortex. *Trends in Neurosciences*, 28, 525–533.
- Hong, S. L., & Newell, K. M. (2006). Practice effects on local and global dynamics of the ski-simulator task. *Experimental Brain Research*, 169, 350–360.
- Ingber, D. E. (2006). Mechanical control of tissue morphogenesis during embryological development. *International Journal of Developmental Biology*, 50, 255–266.
- Kauffman, S. A. (1993). *The origins of order: Self-organization and selection in evolution*. New York: Oxford University Press.
- Keller, E. F. (2002). *Making sense of life*. Cambridge, MA: Harvard University Press.
- Kirkpatrick, S., Gelatt, C. D., & Vecchi, M. P. (1983). Optimization by simulated annealing. *Science*, 220(4598), 671–680.
- Kirschner, M. W., & Gerhart, J. C. (2005). *The plausibility of life*. New Haven, CT: Yale University Press.
- Lungarella, M., & Berthouze, L. (2002). On the interplay between morphological, neural and environmental dynamics. *Adaptive Behavior*, 10(3–4), 223–241.
- Lungarella, M., & Berthouze, L. (2004). Robot bouncing: On the synergy between neural and body–environment dynamics. In F. Iida, R. Pfeifer, L. Steels, & Y. Kuniyoshi (Eds.), *Embodied artificial intelligence* (pp. 86–97). Berlin/Heidelberg: Springer.
- Lungarella, M., Metta, G., Pfeifer R., & Sandini, G. (2003). Developmental robotics: A survey. *Connection Science*, 15, 151–190.
- Mayer-Kress, G., Liu, Y., & Newell, K. M. (2006). Complex systems and human movement. *Complexity*, 12, 40–51.
- Metropolis, N., Rosenbluth, A. W., Rosenbluth, M. N., Teller, A. H., & Teller, E. (1953). Equation of state calculations by fast computing machines. *Journal of Chemical Physics*, 21(6), 1087–1092.
- Meyer, F., Spröwitz, A., & Berthouze, L. (2006). Passive compliance for an RC servo-controlled bouncing robot. *Advanced Robotics*, 20(8), 953–961.
- Platt, R., Fagg, A., & Grupen, R. (2003). Extending fingertip grasping to whole body grasping. *Proceedings of IEEE international conference on robotics and automation*, Taipei.
- Schaal, S., Sternad, D., Osu, R., & Kawato, M. (2004). Rhythmic arm movement is not discrete. *Nature Neuroscience*, 7(10), 1136–1143.
- Sporns, O., & Edelman, G. M. (1993). Solving Bernstein's problem: A proposal for the development of coordinated movement by selection. *Child Development*, 64, 960–981.
- Spröwitz, A., & Berthouze, L. (2005). Robust robot bouncing: Passive compliance and flexible phase locking. *Proceedings of 3rd international symposium on adaptive motion in animals and machines* (p. 59), Ilmenau.
- Thelen, E. (1985). Developmental origins of motor coordination: Leg movements in human infants. *Developmental psychobiology*, 18(1), 1–22.
- Thelen, E. (1994). Three-month-old infants can learn task-specific patterns of interlimb coordination. *Psychological Science*, 5, 280–285.
- Thelen, E., Schöner, G., Scheier, C., & Smith, I. B. (2001). The dynamics of embodiment: A field theory of infant preservative reaching. *Behavioral and Brain Sciences*, 24, 1–86.
- Turing, A. M. (1952). The chemical basis of morphogenesis. *Philosophical Transactions of the Royal Society B (London)*, 237, 37–72.
- Vereijken, B., Whiting, H. T. A., Van Emmerik, R. E. A., & Newell, K. M. (1992). Free(z)ing degrees of freedom in skill acquisition. *Journal of Motor Behavior*, 24, 133–142.
- Warren, W. H. (2006). The dynamics of perception and action. *Psychological Review*, 113, 358–389.

- Webb, B. (2001). Can robots make good models of biological behaviour? *Behavioral and Brain Sciences*, 24, 1033–1050.
- Webb, B. (2002). Robots in invertebrate neuroscience. *Nature*, 417, 359–363.
- Winfrey, A. (1980). *The geometry of biological time*. New York: Springer.
- Wolfram, S. (2002). *A new kind of science*. Champaign, IL: Wolfram Media.
- Zeigler, B. P. (1976). *Theory of modeling and simulation*. New York: Wiley.