Aberystwyth University

Developmental Learning for Autonomous Robots

Meng, Qinggang; Lee, Mark; Chao, Fei

Published in:
Robotics and Autonomous Systems

DOI:
[10.1016/j.robot.2007.05.002](https://doi.org/10.1016/j.robot.2007.05.002)

Publication date:
2007

Citation for published version (APA):
Meng, Q., Lee, M., & Chao, F. (2007). Developmental Learning for Autonomous Robots. *Robotics and Autonomous Systems*, 55(9), 750-759. <https://doi.org/10.1016/j.robot.2007.05.002>

General rights

Copyright and moral rights for the publications made accessible in the Aberystwyth Research Portal (the Institutional Repository) are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the Aberystwyth Research Portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the Aberystwyth Research Portal

Take down policy

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

tel: +44 1970 62 2400
email: is@aber.ac.uk

Developmental Learning for Autonomous Robots

M. H. Lee^{a,*}, Q. Meng^b F. Chao^a

^a*Department of Computer Science, University of Wales, Aberystwyth, SY23 3BD, Wales, UK*

^b*Department of Computer Science, Loughborough University, Loughborough, LE11 3TU, England, UK*

Abstract

Developmental robotics is concerned with the design of algorithms that promote robot adaptation and learning through qualitative growth of behaviour and increasing levels of competence.

This paper uses ideas and inspiration from early infant psychology (up to 3 months of age) to examine how robot systems could discover the structure of their local sensory-motor spaces and learn how to coordinate these for the control of action.

An experimental learning model is described and results from robotic experiments using the model are presented and discussed.

Key words: Sensory-motor coordination, Developmental robotics, Early infant learning

1. Introduction and Approach

Truly autonomous robots have been a major research target throughout the history of Artificial Intelligence, and many avenues have been explored in the search for understanding of how such systems might be created. It is generally agreed that autonomy implies survival over a reasonably long operational period, and this must involve some form of learning because novel experiences will be inevitable and must be handled appropriately. However, the idea that learning abilities might *develop*, rather than being programmed *ab initio*, has not been explored extensively. This situation is beginning to change and the field of developmental robotics is now becoming established as a new strategy for robot learning, see [Lungarella et al., 2003] for a review.

Developmental learning is concerned with not just the ability to learn and gain mastery at a given task but more importantly with how learning may progress and grow to achieve competence over a series of new tasks as they are encountered. This differs from other learning research in that the goal is not concerned with building task specific learning methods but instead aims to find effective algorithms for processes that support the growth of learning, and thus make steps towards the goal of “continuous development”, in the sense of [Prince et al., 2005].

In this paper we describe research into psychologically inspired mechanisms that support developmental behaviour in robotic systems. This approach is fundamentally different from, although complementary to, current work in computational neuroscience. Hence, it is useful to make explicit various aspects of our aims and objectives:

- Computational neuroscience is grounded in structural data and mainly uses the connectionist methods of neural modelling. We can characterise this as a “bottom-up” approach. In contrast, we can describe approaches grounded in behavioural data as “top down”. Rather than build structures that produce desired outputs, we are exploring how complex changes in behaviour patterns could infer appropriate support processes.
- Psychology has generated a rich source of data and theory on human cognitive behaviour but, unlike computational neuroscience, this has not yet been explored to any great extent from a robotics perspective.
- The most dynamic and productive period of human cognitive growth occurs in infancy and the very earliest levels of development are likely to be crucial in determining the structures that later support and build into more advanced skills.
- Consequently, our focus is on basic sensory-motor coordination and control, in particular the emerging control of the limbs and eyes during the first three months of life. In this context the apparently simple act of reaching and grasping an object is a major achievement that relies on prior experience and understanding of sensory spatial encodings, motor properties, and multi-modal coordina-

* corresponding author.

Email address: mhl@aber.ac.uk (M. H. Lee).

tion and control.

- Although the tasks described here can be solved by engineering existing analytical and mathematical methods, such solutions become honed to one task and are generally not extendible to cope with new tasks. This can be seen in the notable difficulties of transferring task experience in connectionist networks [Thrun and Pratt, 1998]. Our interest is in methods that directly deal with this more general issue.
- From this viewpoint some very general requirements can already be discerned. For example, fast, incremental and cumulative mechanisms are desirable; and learning algorithms that need thousands of training cycles can be ruled out.
- Given the above behavioural goals, we do not expect the results to lead immediately to new application solutions. But the findings do shed light on the design of future systems, and these will be important where autonomous systems have to explore unknown environments and need to achieve competencies unknown at design-time.
- The work reported here is part of a larger experimental programme based on hand/eye coordination and control. This paper specifically addresses the spatial control of the arm system.

We emphasise explicit, abstract models and try to avoid preselected internal representations. Our aim is to understand the main design issues, make clear any underlying assumptions, and expose any conditions that support robotic developmental learning.

This paper is structured as follows. Section 2 briefly introduces the psychological background and previous work. Section 3 describes an experimental robot system and the main modelling issues. Section 4 presents results from experiments. The implications of the results are then discussed in Section 5, and Section 6 concludes.

2. Psychological Background and Related Work

Developmental psychology concerns the study of behaviour, and changes in behaviour over time, and attempts to infer internal mechanisms of adaptation that could account for the external manifestations. A key concept in development is the idea of behavioural stages — periods of growth and consolidation — followed by transitions — phases where new behaviour patterns emerge. The idea of stages and staged competence learning was promoted by the prominent psychologist, Jean Piaget [Piaget, 1973]. We believe Piaget’s emphasis on the importance of sensory-motor interaction contains much relevance for robotics. The work of Piaget and others on the plasticity of infant learning has hinted at mechanisms that could underlie stage changes and the growth of competency [Bruner, 1990].

We argue it is necessary to begin investigations at the very earliest levels of development because early experiences and structures are likely to underpin all subsequent growth in ways that may be crucial. This agrees

with the suggestion that sensory-motor coordination is likely to be a significant general principle of cognition [Pfeifer and Scheier, 1997]. Hence, we do not address aspects such as language acquisition but focus on the early sensory-motor learning necessary to control and coordinate a hand/eye system to achieve mastery of reaching and grasping tasks.

The earliest level of sensory-motor development corresponds to the emerging control of the limbs and eyes during the first three months of life. It may seem that the newborn human infant displays no control or coordination but this is a period of enormous growth and adaptation. Many investigators report that early infant life is a period of intense and systematic exploration [Rochat, 2003]. From spontaneous, uncoordinated, apparently random movements of the limbs the infant gradually gains control of the parameters, and coordinates sensory and motor signals to produce purposive acts in egocentric space [Gallahue, 1982]. This is the kind of scenario that has much relevance for autonomous robotics and we need to understand how some of the infant’s learning behaviour might be reproduced.

The most extensive work on computer based Piagetian modeling has been that of Drescher [Drescher, 1991], following early work by Becker [Becker, 1973]. However, Drescher’s system tried to cross-correlate all possible events and was computationally infeasible as a brain model. Maes showed how Drescher’s approach can be improved by using focus of attention mechanisms, specifically using sensory selection and cognitive constraints [Foner and Maes, 1994].

Many models of sensory-motor coordination have been based on connectionist architectures [Kalaska, 1995]. For example, Baraduc et al designed a neural architecture that computes motor commands from arm positions and desired directions [Baraduc et al., 2001]. Other models use basis functions [Pouget and Snyder, 2000] but all these involve weight training schedules that typically require in the region of 20,000 iterations [Baraduc et al., 2001]. They are also tend to use very large numbers of neuronal elements. While the behavioural output of many recent systems has resonances with human behaviour, very few follow the psychological literature on development and even less deal with transitions between qualitatively distinct patterns of behaviour.

3. An Experimental System for Development

A principle central to developmental robotics, and other areas of robotics research, is the concept of embodiment [Clark, 1998]. Effective cognitive agents can not be built as purely symbolic processing engines but must have bodies and physically interact in the world if they are to be capable of learning in the fullest sense [Damasio, 2000, Thelen, 2000].

In our research programme we have built an embodied hand/eye hardware system as a test-bed for exploring and assessing different developmental algorithms. This labora-

tory robot consists of two industrial quality manipulator arms and a motorised camera system. These are configured in a manner similar to the spatial arrangement of an infant's arms and head — the arms are mounted, spaced apart, on a vertical backplane and operate in the horizontal plane, working a few centimeters above a work surface, while a computer-controlled pan and tilt colour camera, is mounted above and looks down on the work area. Figure 1 shows a view of one of the robot arms.

Fig. 1. The experimental robot arm system

Each rotational joint of the robot arms has a motor drive and also an encoder which senses the joint angle, thus providing a proprioceptive sense. The arms carry a very simple tactile sensor in the form of a fingertip contact device at the limb end-point. This fingertip probe consists of a 10mm diameter rod containing a small proximity sensor. The sensor faces downwards so that as the arm sweeps across the work surface any objects passed underneath will be detected. Normally, small objects will not be disturbed but if an object is taller than the arm/table gap then it may be swept out of the environment during arm action.

The camera system is arranged to fixate and perform saccades like an eye. However, the experiments described here deal with very early motor learning before grasping behaviour has developed and before visually guided action has emerged, hence the visual modality is not discussed further in this paper.

To illustrate the limb geometry, figure 2 shows a diagram of the configuration of an arm. The two degrees of freedom can be specified by the joint angles, θ_1 and θ_2 .

The task for this sensory-motor system is to learn the effects and operational properties of its various motors and sensors, and then learn to control and coordinate them to achieve reaching behaviour towards stimulating objects. We can now see the problem in raw form: given an embodiment (or anatomy), an environment and a set of basic sensory-motor functions, how can we best discover their properties and learn to use them in a framework of developing competence? By considering the logical possibilities for this simple sensory-motor system we see various fundamental questions:

Sensory encoding — What coding scheme is appropriate or effective for spatial sensory feedback? This is important as it will have consequences for any representa-

Fig. 2. A plan view of the arm spatial configuration.

tion of local space and will thus affect how space can be learned and used.

Motor control — How can new actions develop? Clearly, motor action can occur in many forms but what is a suitable motor behaviour for learning in this situation?

Coordination — In this intra-modal study we can explore the coordination of motor action with proprioception and touch sensing. The essential correlation between sensor spaces and motor spaces seems to be a foundation stone for development, and occurs at many levels [Pfeifer and Scheier, 1997].

Constraint schedules — If behavioural stages are shaped by constraints, then our understanding of stages will be increased by exploring how constraints could or should be exploited.

Motivation — Some driving force is necessary to cause unsupervised learning to occur. We find a simple novelty function is sufficient to motivate all action.

The following subsections give more detail on these modelling and design issues.

3.1. An abstract motor model

We use a very abstract motor model for the movement of the robot arms described in section 3. As seen in figure 2 there are two independent degrees-of-freedom, both based on revolute joints. It is not necessary to model at the level of extensors and flexor muscle pairs and we use two variables, M_1 and M_2 which represent the overall drive strength applied to the respective joint. The force applied to a limb section can be opposed by various effects but we assume that viscous properties from the many biochemical processes operating during skeletal muscle contraction are dominant [Winter, 1990]. Arm mass is also important but is often balanced out within the internal force interplay between flexors and extensors, and elasticity, usually from the tendons, tends to be significant only at the far extent of movement. This means, to a first approximation, the forces given by M_i will determine the speeds of the limb segments being driven, i.e. $M_i = k_2 \dot{\theta}$.

By integrating each M_i over time during an action we can obtain a set of values which represent the “extent” of an action for each degree-of-freedom: $d_i = \int M_i dt$. Thus, d_i is a measure of the extent of movement of limb segment i during a time interval under motor value M_i . An integrator is assigned to each motor system to generate these extent values. While more detailed models can be advocated, an extent signal can always be extracted from the motor formulation and this simplified model seems to be a reasonable abstraction for many practical purposes [Winter, 1990].

In our experiments we operate the arms at a slow rate in order to minimise the effects of external forces generated by dynamics and gravity. In more dynamic situations the ill-posed problems of motor control become important and require complex compensatory feedback [Bernstein, 1967].

3.2. Proprioception and tactile sensing

By virtue of their differing physical structure and constraints, each modality will have its own coding of space. Thus, when an eye refers to a spatial location then that data will only have meaning in terms of the actions required to move or direct the eye to that position. Similarly for a limb; for example, locations in proprioception space are encodings of signals that correspond to the arm being at a certain location.

The main sensing systems in human limbs include the mechanoreceptors consisting of the internal proprioception sensors and the surface-based tactile or contact sensors. Proprioception provides feedback on the sensed position of the limb in space. Although the actual biological mechanisms of proprioceptive feedback are not entirely known, we do know that the angles between limbs at joints are sensed and other more complex kinaesthetic responses can be found [Bosco et al., 2000]. The question arises as to whether any particular form of neural encoding of limb position has more functional efficacy than others.

From figure 2 there are two joint angles, θ_1 and θ_2 , i.e. two degrees-of-freedom, and so we need two state variables; let these be S_1 and S_2 . The proprioception encoding question then concerns what could or should these variables represent.

The simplest encoding scheme is for the proprioceptive neurons to generate feedback signals based directly on the angles of displacement at the joints. Thus,

$$S_1 = f(\theta_1) \quad S_2 = f(\theta_2)$$

where f is a linear or at least smooth monotonic function. We refer to this encoding as a **joint angle** coordinate scheme. Such feedback is known to be produced by the afferent fibres from mechanoreceptors embedded in the joint capsules [Bosco et al., 2000].

Note that if the spatial location of the limb end-point could be sensed then this would lead to a much more useful encoding of space as most actions are concerned with driving the end-point (the hand, or end-effector) to desired

spatial locations, thus any sensory feedback signals would better relate to the control of actions. However, this makes the production of the sensory signals much more complex:

$$S_1 = \sqrt{l_1^2 + l_2^2 + 2l_1l_2 \cos \theta_2}$$

$$S_2 = \theta_1 - \arctan \frac{l_2 \sin \theta_2}{l_1 + l_2 \cos \theta_2}$$

where l_1 and l_2 are the lengths of the upper-arm and forearm respectively. Here S_1 is the effective length of the arm axis from shoulder to hand and S_2 is the angle this axis makes with the baseline at the shoulder. We refer to this coordinate frame as a **shoulder** encoding.

Another, even more attractive scheme, would be to relate the arm end-points to the body centre-line. This **body-centred** encoding would be appropriate for an ego-centric space as might be used in very early development. To obtain a **body-centred** encoding we modify the shoulder encoding by shifting the reference point from the shoulder to the centre of the body, (i.e. the distance *Base*). Then S_1 is the distance from the body to the hand and S_2 is the angle between the hand and the body centre-line, thus:

$$S_1 = \sqrt{(Base - \bar{X})^2 + \bar{Y}^2}$$

$$S_2 = \arctan \frac{\bar{Y}}{Base - \bar{X}}$$

$$\begin{aligned} \text{where } \bar{X} &= (l_1 \cos \theta_1 + l_2 \cos(\theta_1 - \theta_2)) \\ \text{and } \bar{Y} &= (l_1 \sin \theta_1 + l_2 \sin(\theta_1 - \theta_2)) \end{aligned}$$

One other notable spatial encoding is a frame where the orthogonal coordinates are lateral distance (left and right) and distance from the body (near and far). The signals for this case are the coordinate values of the end-points in a rectangular space, thus:

$$S_1 = l_1 \cos \theta_1 + l_2 \cos(\theta_1 - \theta_2)$$

$$S_2 = l_1 \sin \theta_1 + l_2 \sin(\theta_1 - \theta_2)$$

This encoding, referred to as **Cartesian** encoding, seems the most unlikely for a biological system, as it is the most abstract being based on a rectilinear rather than polar coordinate frame and has no reference point to the body. However we include this scheme due to its ubiquity and apparent importance in human spatial reasoning [Newcombe and Huttenlocher, 2000].

Before vision comes into play, it is difficult to see how such complex feedback as given by the three latter encodings could be generated and calibrated for local space. The dependency on trigonometrical relations and limb lengths at a time when the limbs are growing significantly makes it unlikely that these codings could be phylogenetically evolved. Only the **joint angle** scheme could be effective immediately but the others may develop through growth processes. Recent research [Bosco et al., 2000] on the hind limbs of adult cats has discovered that *both joint angle and shoulder* encodings can coexist, with some neuronal groups giving joint angle outputs while others give foot/hand position

encodings independently of limb geometry. We investigated all four systems as candidate encodings for proprioception signals.

3.3. *Motivation and Reflex Action*

A motivational component is necessary to drive learning and there is evidence from infant studies that novelty is a strong motivational stimulus. Our parsimonious approach dictates that we look for a very simple novelty function and we define any new sensory value change that has not been experienced recently as a novel stimulus. The values of prior stimuli are recorded in the sensory-motor maps and any incoming sensed values can be compared with these “expected” values. Any new or changed stimulus is assigned a high field excitation level and the system attends to the field with the highest excitation. Habituation mechanisms are used to reduce excitation with repetition and time and so attention is attracted by novelty and decays with familiarity. This motivational mechanism automatically attempts to repeat actions that accompanied the most exciting stimuli until eventually attention shifts away when familiarity has been gained.

However, without any initial stimulus or prior experience there would be no reason to act and hence we must provide a basic “reflex” to initiate the system. As there are no grounds for selecting any particular motor values the choice is between random values or some preset reflex action. From the literature, we favour an energetic burst of activity as a reflex act. Notice that when the arms are in the lateral position there is less opportunity for interaction or stimulation and this forms a natural “rest” position. This position also has less kinematic freedom (see figure 2) and consequently there is less positional ambiguity when the arm is returned to this location. Furthermore, if the hands are in the lateral position and the motor values are set to a high positive value then this gives a general sweep movement. This reflex action of driving the motors “full on” brings the hand from the rest position to the body centre-line (in a position roughly equivalent to the mouth). The arm is returned to the rest position after any action and so the start location is consistent. This gives a reflexive homing behaviour and the rest area provides a kind of fiducial point for the execution of actions. The drive position integrators are reset to zero whenever the arm reaches the rest area.

The ballistic approach to motor action is readily seen in leg action in three month old infants. In motor experiments where kicking behaviour is able to disturb a stimulus, infants learn to adapt their kicking to achieve a desired stimulus change [Thelen and Fisher, 1983]. Other work has shown that 3 month old infants have considerable proprioceptive control of their legs and without any visual input can learn different action patterns [Angulo-Kinzler et al., 2002].

3.4. *Constraint Lifting*

Human cognitive development has been characterised by progression through distinct stages of competence, each stage building on accumulated experience from the level before. This can be achieved by lifting constraints when high competence at a level has been reached [Rutkowska, 1994]. Any constraint on sensing or action effectively reduces the complexity of the inputs and/or possible action, thus reducing the task space and providing a frame or scaffold which shapes learning [Bruner, 1990, Rutkowska, 1994]. Such constraints have been observed or postulated in the form of sensory restrictions, environmental constraints, anatomical limitations, and internal cognitive or computational limits [Hendriks-Jensen, 1996]. Internal sensory and motor constraints are evident in the newborn, for example the visual field begins as a kind of tunnel vision [Hainline, 1998] and the width of view has been reported as growing from 30 degrees at 2 weeks of age to 60 degrees at 10 weeks [Tronick, 1972].

We view “constraint lifting” as a key mechanism for progression towards increasing competence. Transitions between stages are related to the lifting of constraints, although the nature of such transitions is not fully understood. It is clear that transitions must be related to internal global states, not local events, and we obtain certain global values by simple summation over internal local properties. For example, the global sum of stimuli excitation levels will reach a plateau when no novel events have been seen for some time. Such global values can signal the need to enter a new level of learning by lifting a constraint or accessing a new sensory input.

3.5. *A Mapping Model for Sensory-Motor Coordination and Learning*

Any model of sensory-motor learning will need some form of computational substrate or representation upon which the dynamic relationships between sensory and motor events are experienced, learned and explored. We have developed such a computational substrate as a framework for investigation and for the development of algorithms. Our model is based on a two-dimensional mapping scheme.

All the mappings used in this work consist of two-dimensional sheets of elements, each element being represented by a patch of receptive area known as a **field**. The fields are circular, regularly spaced, and overlapping. Only two parameters are needed to define a uniform map structure: field size and inter-field spacing. These determine the degree of field overlap and the field density, per surface unit area. Every field in a map has a set of associated variables that can record state information, this includes sensory stimulus values and excitation levels as mentioned in section 3.4. The field values in a map decay with time and so the maps can be viewed as a form of short term memory.

Each map layer deals with a single channel of either sen-

sory or motor information and corresponding fields between layers are directly linked. Figure 3 shows two map layers, the upper layer is a kinaesthetic sensory map, in Cartesian encoding, and the lower layer is the associated motor drive map. Only boundary fields are shown for clarity (the internal areas become filled through tactile contact experience). Note that each point in one map will have an associated field in the other. The labels in figure 3 indicate the correlation between the maps and significant mutual distortion or warping is clearly visible.

Fig. 3. Map correlation. Each field in the sensory map (upper) is connected to a field in the motor map (lower). The shaded fields, A, B, C, D, indicate the associations between the maps.

There are various possible ways of creating fields and we have experimented with methods for growing fields of various sizes and locations on demand [Meng and Lee, 2007]. However, in this study we use uniform sheets of identical field sizes and spacings. These basic uniform maps are not pre-wired or pre-structured for any specific spatial system and fields are created by simply assigning them to new sensory or motor values. We use two access variables, X, Y , to reference locations on any given map; these simply define a point on the two-dimensional surface — it is important to state that they do not have any *intrinsic* relation with any external space. Thus, a map starts as an empty sheet, and the fields, when addressed through the access variables, become populated with sensory or motor data for experiential events.

4. Experiments and results

The experiments were designed to probe the research questions given in Section 3. We experimented with active and inactive contact sensing as this is an internal constraint that clearly affects behaviour. Other constraints such as

sensory resolution and motor patterns were also explored directly.

The effects of different field sizes were examined by creating three maps, each of different density. The different field sizes used, termed small, medium and large, were of diameter 10, 20 and 40 units, respectively. The incoming S and M signals were used to build each map separately and simultaneously. However, attention and action selection can only choose one field, so we drive actions from each map in turn in separate experiments and can then observe the behaviour and effectiveness of the mapping parameters.

For proprioception there are four candidate encodings of the proprioceptive signals, S_1 and S_2 , as described in Section 3.2. Each of these were tried in turn to determine their effect on the shape of the developing sensory-motor space.

4.1. Results

The first trials began with no contact sensing and no prior experience. Any objects were either ignored or pushed out of range. Figure 4 illustrates the behaviour as traces of movements — for clarity these are displayed as directed lines between start and end points in motor space. From

Fig. 4. Arm movements with no contact sensing. Actions are shown as vectors in motor space.

this figure we see that the arm first moved repeatedly between the rest area (lower right) and the body area (centre left). But as the stimulation for this action habituated so global excitation levels fell and spontaneous moves were introduced, leading to fields on the boundary being discovered and explored. Figure 5 shows the fields discovered after the above trial — this diagram is in Cartesian space to show the locations in relation to the physical arm geometry.

Eventually a plateau in field growth was reached and this was used as the trigger to lift a constraint, in this case by enabling contact sensing. Figure 6 shows rest/body moves being interrupted by contact with an object on the hand path. Such contact events create internal (non-boundary) fields, as seen in figure 7, and generate much repeated action.

The pattern of new field growth is instructive and figure 8 illustrates this with a plot of the total number of new fields produced over time. The fields have been classified into two types: boundary fields and internal (contact) fields, and

Fig. 5. Fields generated during non-contact stage. Shown in Cartesian encoding of proprioception space.

Fig. 6. Arm movements with active contact sensing. An object (near the centre of the diagram) caused sensory interrupts which were followed by repeated contact action.

Fig. 7. Fields generated after first object contact. Shown in Cartesian encoding.

their numbers plotted against trials. Initially there were no objects in the environment and so only boundary fields are discovered; 54 fields had been created by trial 100, this is around 85% of the total possible on the boundary for this map. Then two objects were introduced at trial 106 and a growth in internal fields begins. At trial 180 the two objects' positions were altered and further fields were then rapidly created. During this period no more boundary fields are discovered. Eventually, no more internal fields can be

produced by the presence of the objects and from around trials 300 to 400 some spontaneous action finds a few more boundary fields. At trial 400 two more objects are introduced and further internal field activity occurs. The characteristic plateau shape is seen to emerge for the boundary fields, and this eventually occurs for the internal fields too.

Fig. 8. Rates of growth in maps. Only initial field visits are counted.

From the experiments we observe a progression of qualitatively distinct behaviour patterns: (1) “blind groping” actions mainly directed at the body area, (2) more groping but at the boundary regions, (3) unaware pushing of objects out of the local environment, (3) limb movements stopping upon sensing object contact, (4) repeated cycles of contact and movement, i.e. repeated “touching” of detected objects. (5) directed touching of objects and sequences of objects. In the last case, if objects exist at several locations then attention will shift to each object in turn, as they alternatively become habituated and stimulated, so that a roughly cyclic behaviour pattern is produced, similar to eye scanpaths, see Figure 9. All these behaviours, including

Fig. 9. Cyclic patterns of behaviour produced by several objects.

motor babbling and the rather ballistic motor actions, are widely reported in young infants [Piek and Carman, 1994].

5. Discussion

There were four options for the encoding of proprioception signals and, figure 10 shows the four maps produced from the experiments. The encodings were: Joint, Shoulder, Body and Cartesian, and the marked fields and the arrows indicate the correspondences between the maps. For example, moving through the fields in the direction of arrow 1

Fig. 10. Maps for the 4 proprioceptive encodings. From the left: Joint, Shoulder, Body and Cartesian encodings.

will identify equivalent fields in each map. Only boundary fields are shown in order to illustrate the map shapes and correspondences more clearly.

Although there are various relative distortions, all the schemes are continuous, smooth non-linear mappings and so they simply represent different distortions or warpings of the two-dimensional sheet. At first glance there does not seem to be any clear advantage for any one of the four encoding schemes, as all can successfully support the coordination function. However, we can make some interesting observations.

As mentioned in Section 3.2, although the other encodings are more desirable, only the joint signals could be expected to be available in very early life. However, when looking for the biological sources of proprioception we do not find many joint angle sensors, although they do exist. Instead, there are many muscle spindle receptors but these detect linear stretch in the muscles, not rotation [Prud'homme and Kalaska, 1994]. However, it turns out that the spindle signals are even more appropriate for the spatial encodings. If we assume that two jointed limb segments are 20 units long and a muscle is affixed at 1 unit from the joint on one limb and 20 units from the joint on the other, then the length of the muscle will be: $\sqrt{401 + 40 \cos \theta}$. When this is plotted against the shoulder-hand distance as the joint rotates through 180 degrees it becomes clear that the spindle gives a more linear approximation to the required distance measure than does the joint angle, see Figure 11. In particular, the spindle signal is most linear in the region where the joint angle is least useful (tending to infinite slope). The figure also shows a combined curve that is a composite of the two signals. This suggests that the shoulder encoding may not be too difficult to implement biologically and hence, the other trigonometric functions may also be possible. Certainly, muscle stretch signals may be at least as important as joint angle receptors, and this analysis shows that there are clear advantages if they coexist together. This is supported by recent neurological findings [Bosco et al., 2000]. Other related research [Bosco and Poppele, 2000] has indicated that many of the neurons involved in proprioception are not affected by muscle forces or other dynamics and hence act as pure kinematic sensors.

To complete this analysis, we can also compute the effect

Fig. 11. Sensing the shoulder-hand distance. $\square$ gives joint angle in degrees, $\triangle$ is muscle spindle (scaled), and $\circ$ is a composite of both curves.

of using spindle sensors for the body-centred encoding. For comparison figure 12 shows the body-hand distance computed for both the joint angle ranges used in the experiments and computed for two muscle spindle sensors over the equivalent ranges. The stretch sensors continue to give the more linear relationship and so this way of computing the kinematics may be both biologically plausible (and better than sensing solely joint angles) as well as an effective way of measuring important variables without having to perform complex trigonometrical transformations.

Another aspect of proprioception is that the joint angle and shoulder encodings are both **local** to their individual limb, whereas the body-centred and Cartesian encodings can both provide a **common** framework that includes both limbs. The shoulder encoding has a separate origin for each arm but the body encoding uses the same framework for *both* arms. This suggests the body-centre has an important role as a key fiducial point in egocentric space. Indeed, any head-mounted sensing system, such as the eyes, will also have a polar, body-centred reference frame which would naturally align with a body-centred proprioception system.

The Cartesian scheme has no natural anchor point but if we set the origin at the body centre then, at a reasonable distance, the body-centred encoding approximates to a rectangular grid that can matches a regular Cartesian scheme. For example, the errors between these two schemes increase nearer the body and with increasing angular displacement from the centre-line. However, there is an operating region where a close match occurs and this may be

Fig. 12. Body-hand distance plotted against joint angles (upper) and muscle stretch sensors (lower).

relevant for the genesis and exploitation of the Cartesian frame in human space. In fact, the close relationship between body-centred space and Cartesian space can be seen in figure 10 where the curved pattern is very similar for the two relevant maps and are particularly well matched in the central region. This suggests how allocentric spatial frames may eventually emerge from the more basic but ultimately limited egocentric frames.

For the effects of map field sizes we find a trade off, between speed of exploration and accuracy of motor acts. When larger fields are used they cover more sensory space and thus the full mapping is learned much faster. However, larger fields generalise many sensory signals into one spatial representation. If smaller fields are used then the specification of sensory space is more acute and movements to given locations are more likely to be accurate, but much more exploration is needed to generate the mappings. We found it best to start with the coarsest map and then, when all the fields had been accessed (as signaled by a global variable), the system transitioned to a finer scale map. This effect was also reported by [Gomez, 2004]. Figure 13 shows this progressive transition over the 3 map sizes; the switches occurred at trials 34 and 68. It is interesting that the receptive field size of visual neurons in infants is reported to decrease with age and development and this leads to more selective responses [Westermann and Mareschal, 2004].

One of the most central issues in developmental studies concerns the organisation of behavioural development: that is, are there any preferred schedules for constraint-lifting or reflex suppression that are best for learning and growth? It has been long believed that visually guided reaching is the earliest accurate reaching behaviour to occur. Infants spend time observing their hands around 12 weeks and “visually guided” reaching begins between 15 and 20 weeks. Reaching after 22 weeks is visually triggered rather than

Fig. 13. Transitions between three maps of different scale. Only initial field visits are counted. The Mapping Scale plot indicates the switching points.

guided. However, [Clifton et al., 1993] have performed infant reaching experiments in the dark and shown that infants are able to use proprioception, not vision, in successful reaching tasks. A form of “hand looking” behaviour is bound to occur when the hand first enters the visual field as an “unknown” object; but the question is whether this stage is essential to, and therefore must occur before, visually-guided behaviour or whether there could be other schedules. Our study confirms the view of Clifton *et al* by showing how proprioceptive learning can guide action, can be prior to visual development and does not depend upon visual confirmation. Our model will be well placed to support the next stages of hand-looking and visual-guidance but these should be faster and more robust by drawing on a well developed kinaesthetic underpinning of local space. As Clifton *et al* state: “Prior accounts of early reaching have underemphasized the role of proprioception in infants’ acquisition of prehension” [Clifton et al., 1993].

6. Conclusions

The system described in this paper constructs sensory-motor schemas in terms of interlinked topological mappings of sensory-motor events, pays attention to novel or recent stimuli, repeats successful behaviour, and detects when reasonable competence at a level has been achieved. The behaviour observed from the experiments displays an increasing progression from initially spontaneous limb movements, followed by more exploratory movements, and then directed action towards touching and tracing objects. During these stages the local egocentric limb spaces become developed and form a substrate for future cross-modal skilled behaviours.

Our method is incremental and fast, using each experience wherever possible. This is in accord with several researchers who report that learning and adaptation can be very fast [Angulo-Kinzler et al., 2002, Rochat and Striano, 1999] and often seem to need only one trial to alter behaviour. We believe this is an important feature for autonomous systems in real environments and applications.

Our approach has been supported by the findings cited and reports such as [Gomez, 2004] who show that starting with low resolution in sensors and motor systems, and then increasing resolution, leads to more effective learning. The reduction in degrees of freedom obtained by staged development is also reported to be an effective strategy [Lungarella and Berthouze, 2002], as is the concept of constraints being beneficial to the emergence of stable patterns and helping to bootstrap later stages [Berthouze and Lungarella, 2004]. As an early researcher stated:

“Gradual removal of constraint could account for qualitative change in behaviour without structural change” [Tronick, 1972]

Acknowledgements

We thank EPSRC for grant GR/R69679/01.

References

- [Angulo-Kinzler et al., 2002] Angulo-Kinzler, R., Ulrich, B., and Thelen, E. (2002). Three-month-old infants can select specific leg motor solutions. *Motor Control*, 6(1):52–68.
- [Baraduc et al., 2001] Baraduc, P., Guigon, E., and Burnod, Y. (2001). Recoding arm position to learn visuomotor transformations. *Cerebral Cortex*, 11:906–917.
- [Becker, 1973] Becker, J. D. (1973). A model for the encoding of experimental information. In Schank, R. C. and Colby, K. M., (Eds.), *Computer models of thought and language*, pages 396–434. W. H. Freeman and Company, San Francisco, USA.
- [Bernstein, 1967] Bernstein, N. (1967). *The Coordination and Regulation of Movement*. Pergamon Press, Oxford.
- [Berthouze and Lungarella, 2004] Berthouze, L. and 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):47–64.
- [Bosco and Poppele, 2000] Bosco, G. and Poppele, R. (2000). Reference Frames for Spinal Proprioception: Kinematics Based or Kinetics Based? *J. Neurophysiol.*, 83(5):2946–2955.
- [Bosco et al., 2000] Bosco, G., Poppele, R., and Eian, J. (2000). Reference frames for spinal proprioception: Limb endpoint based or joint-level based? *J. Neurophysiol.*, 83(5):2931–45.
- [Bruner, 1990] Bruner, J. (1990). *Acts of Meaning*. Harvard University Press, Cambridge, MA.
- [Clark, 1998] Clark, A. (1998). *Being There — Putting Brain, Body, and World Together Again*. MIT Press, Cambridge, MA.
- [Clifton et al., 1993] Clifton, R., Muir, D., Ashmead, D., and Clarkson, M. (1993). Is visually guided reaching in early infancy a myth? *Child Development*, 64(4):1099 – 1110.
- [Damasio, 2000] Damasio, A. (2000). *The Feeling of what Happens: Body and Emotion in the Making of Consciousness*. Vintage.
- [Drescher, 1991] Drescher, G. (1991). *Made up minds: a constructivist approach to artificial intelligence*. MIT Press, Cambridge, MA.
- [Foner and Maes, 1994] Foner, L. and Maes, P. (1994). Paying attention to what’s important: Using focus of attention to improve unsupervised learning. In *Proc. 3rd Int. Conf. Simulation of Adaptive Behaviour*.
- [Gallahue, 1982] Gallahue, D. (1982). *Understanding Motor Development in Children*. John Wiley, NY.
- [Gomez, 2004] Gomez, G. (2004). Simulating development in a real robot. In *Proceedings of the 4th Int. Workshop on Epigenetic Robotics*, pages 1–24.
- [Hainline, 1998] Hainline, L. (1998). How the visual system develops. In Slater, A., (Ed.), *Perceptual development: Visual, auditory, and speech perception in infancy*, pages 5–50. Psychology Press.
- [Hendriks-Jensen, 1996] Hendriks-Jensen, H. (1996). *Catching Ourselves in the Act*. MIT Press, Cambridge, MA.
- [Kalaska, 1995] Kalaska, J. (1995). Reaching movements: Implications of connectionist models. In Arbib, M. A., (Ed.), *The Handbook of Brain Theory and Neural Networks*, pages 788–793. MIT Press.
- [Lungarella and Berthouze, 2002] Lungarella, M. and Berthouze, L. (2002). Adaptivity through physical immaturity. In *Proceedings of the Second International Workshop on Epigenetic Robotics*, volume 94, pages 79–86. Lund University cognitive studies.
- [Lungarella et al., 2003] Lungarella, M., Metta, G., Pfeifer, R., and Sandini, G. (2003). Developmental robotics: a survey. *Connection Science*, 15(4):151–190.
- [Meng and Lee, 2007] Meng, Q. and Lee, M. H. (2007). Automated cross-modal mapping in robotic eye/hand systems using plastic radial basis function networks. *Connection Science*. preprint available from the authors.
- [Newcombe and Huttenlocher, 2000] Newcombe, N. and Huttenlocher, J. (2000). *Making Space*. MIT Press, Cambridge, MA.
- [Pfeifer and Scheier, 1997] Pfeifer, R. and Scheier, C. (1997). Sensory-motor coordination: the metaphor and beyond. *Robotics and Autonomous Systems*, 20(2):157–178.
- [Piaget, 1973] Piaget, J. (1973). *The Child’s Conception of the World*. Paladin, London.
- [Piek and Carman, 1994] Piek, J. and Carman, R. (1994). Developmental profiles of spontaneous movements in infants. *Early Human Development*, 39:109–126.
- [Pouget and Snyder, 2000] Pouget, A. and Snyder, L. (2000). Computational approaches to sensorimotor transformations. *Nature neuroscience*, 3:1192–1198.
- [Prince et al., 2005] Prince, C., Helder, N., and Hollich, G. (2005). Ongoing Emergence: A Core Concept in Epigenetic Robotics. In *Proceedings of the 5th Int. Workshop on Epigenetic Robotics*, pages 63–70. Lund University cognitive studies.
- [Prud’homme and Kalaska, 1994] Prud’homme, M. and Kalaska, J. (1994). Proprioceptive activity in primate primary somatosensory cortex during active arm reaching movements. *Journal of Neurophysiology*, 72(5):2280–2301.
- [Rochat, 2003] Rochat, P. (2003). Five levels of self-awareness as they unfold early in life. *Consciousness and Cognition*, 12(4):717–731.
- [Rochat and Striano, 1999] Rochat, P. and Striano, T. (1999). Emerging self-exploration by 2-month-old infants. *Developmental Science*, 2(2):206–218.
- [Rutkowska, 1994] Rutkowska, J. (1994). Scaling up sensorimotor systems: Constraints from human infancy. *Adaptive Behaviour*, 2:349–373.
- [Thelen, 2000] Thelen, E. (2000). Grounded in the world: Developmental origins of the embodied mind. *Infancy*, 1(1):3–28.
- [Thelen and Fisher, 1983] Thelen, E. and Fisher, D. (1983). The organization of spontaneous leg movements in newborn infants. *Motor Behaviour*, 15:353 – 377.
- [Thrun and Pratt, 1998] Thrun, S. and Pratt, L. (1998). *Learning to Learn*. Kluwer Academic Pub.
- [Tronick, 1972] Tronick, E. (1972). Stimulus control and the growth of the infant’s effective visual field. *Perception and Psychophysics*, 11(5):373 – 376.
- [Westermann and Mareschal, 2004] Westermann, G. and Mareschal, D. (2004). From parts to wholes: Mechanisms of development in infant visual object processing. *Infancy*, 5(2):131–151.
- [Winter, 1990] Winter, D. (1990). *Biomechanics and motor control of human movement*. Wiley New York.