Aberystwyth University

A Content-Neutral Approach for Sensory-Motor Learning in Developmental Robotics
Lee, Mark; Meng, Qinggang; Chao, Fei

Published in:
Sixth International Conference on Epigenetic Robotics

Publication date:
2006

Citation for published version (APA):
Lee, M., Meng, Q., & Chao, F. (2006). A Content-Neutral Approach for Sensory-Motor Learning in Developmental Robotics. In *Sixth International Conference on Epigenetic Robotics* (pp. 55-62)
http://hdl.handle.net/2160/471

A Content-Neutral Approach for Sensory-Motor Learning in Developmental Robotics

M. H. Lee   Q. Meng   F. Chao

Department of Computer Science
University of Wales, Aberystwyth, UK
mhl@aber.ac.uk

Abstract

Many computational models of development show how learning may be implemented but they do not often address the relative criticality of the various supporting elements. This paper draws on experience from experimental work in developmental robotics to illustrate how psychological data may drive a content-neutral approach as a complementary alternative to bottom-up neurological models.

1. Introduction

Developmental psychology appears to offer exactly the kind of ideas that are needed for epigenetic robotics. However there is a huge gap between psychological theories and working algorithms for robots. This paper argues for a particular approach towards bridging the gap and illustrates with current experimental work.

Although epigenetic robotics is a new and recent focus for research (Lungarella et al., 2003), there have been many past experiments under the banner of Robot Learning and most avenues and techniques in Artificial Intelligence have been explored in one form or another. However, nearly all such work does not address or even consider developmental processes and frequently, a bias towards some particular mechanism or technique tends to be in evidence, often obscuring any developmental possibilities. For example, the recent popularity of artificial neural network models has often produced studies where the key issue might be phrased as "can a neural net of type X effectively model the growth of skill Y ?". Often the result is successful but, in our view, this usually tells us less about how to implement or understand skill Y and more about the wider generality or applicability of technique X. The question we believe should be asked is "what are the main requirements or features of a mechanism that could produce increasing competence or growth in skill Y ?".

Another issue relates to the debate between the psychological empiricists and nativists; a debate which is still unresolved after fifty years (Thelen and Whitmyer, 2005). This concern with whether certain "core concepts" are learned through sensory-motor behaviour or are innate, has highlighted the significance of mental representations. If strong representations are built into a system then it will be open to the criticism that prior knowledge (concepts or beliefs) is being used before its requirements and origins have been established. We agree with (Thelen and Whitmyer, 2005) that it is important to avoid this by adopting a "content neutral" methodology, i.e. avoiding any ascription of beliefs or intentions. We also agree that representations can not exist independently from the real-time processes, events and environments which define the agent. Consequently, our approach has parallels with that of Thelen *et al* but with the key exception that we do not use continuous time dynamics. General principles for developmental robotics have been described before, e.g. (Prince et al., 2005), but it is useful to list some principles for our approach as follows:

Open methods. The requirements for any particular method should emerge from the properties of the task, and we should aim for (a) minimal representations, and (b) avoid any bias towards pre-selected mechanisms.

Content-neutral. Mechanisms or models should aim for generality of structure, and avoid assumptions about beliefs or other internal cognitive states.

Synthesis plus analysis. The synthesis of new mechanisms can be very productive and can offer complementary insights compared with traditional analysis. This originates in Braitenberg's approach towards understanding complex neural problems (Braitenberg, 1984).

Abstraction. In contrast to bottom-up methods (e.g. neural modelling), top-down models can help to manage complexity and assist understanding. This implies a high level of abstraction and it may be better to start from a very high level and then refine as necessary, rather than

carry forward much complex detail and face later difficulties of interpretation and analysis.

Parsimony. Simplicity is extremely valuable in aiding understanding and interpreting results. The mechanisms or models we wish to find should not have any more complexity than is absolutely necessary. Parsimony and abstraction are valuable tools in facilitating *repeatability* of experiments and investigations.

Behaviour-based. Development is essentially grounded in behaviour, and is observed, detected and measured in terms of behaviour. Thus, results need to be assessed in terms of qualitative and quantitative measures of behaviour.

Data from psychology. Psychologists are the professionals who study behaviour, and it is developmental psychology, both theoretical and experimental, that provides the source of data on behaviour. This rather obvious point often seems to be overlooked in computational and artificial systems work.

Early stages first. In development, natural or artificial, each stage of growth logically builds on and depends upon the preceding stages. This argues for obtaining a good understanding of the earliest stages first, and we believe that there is a great deal to be learned from the youngest infants that will be very significant for studies on later cognitive stages.

Constraints are important. Constraints can exist in many forms: physical (morphology, mechanical, motor); internal (cognitive, sensory, neural, maturational); and environmental (external, scaffolding, social). Such constraints can be very beneficial in reducing the complexity of a task, and actually may be a major controlling factor in development. We believe constraint-lifting could be a key mechanism for developmental algorithms.

2. A Sensory-Motor Development Problem

An experimental system is now presented through which we illustrate our approach. It is important to note that, because of lack of space, this paper is descriptive and does not contain all the mathematical fine detail. The goal here is to explain and discuss the methodological approach.

We are involved in a research programme that is exploring the design of algorithms that allow a robotic hand/eye system to learn to coordinate its sensory-motor apparatus with the goal of producing skilled reaching behaviour. Such visually initiated reaching involves the visual detection of targets, using both peripheral and fixated foveal vision, cross-modal spatial transfer of visual and kinaesthetic information, motor and proprioceptive coordination, and the integration of tactile sensing.

We note that it is necessary to use hardware in robotics because a set of rich visual, tactile and motor spaces are essential for delivering the unexpected artefacts and difficulties posed by the real, noisy world. This also relates to the embodiment principle whereby 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 (Clark, 1998, Thelen, 2000). Figure 1 shows a view of our experimental laboratory robot system. The system consists of two manipulator arms and a
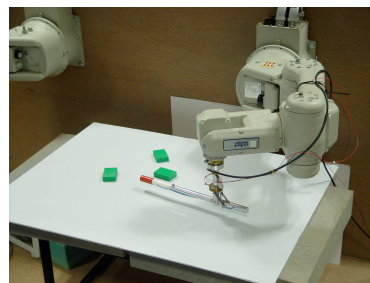

Figure 1: Part of the experimental robot system

vision system that acts as an "eye". 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 centimetres above a work surface, while the "eye", which is a computer-controlled pan and tilt head with a colour imaging camera, is mounted above and looks down on the work area.

In considering the development of skilled reaching we note that many computational models of reaching, e.g. (Baraduc et al., 2001, Ajemian et al., 2000, Mussa-Ivaldi and Giszter, 1992) deal with the learning of a particular task but do not address the relative criticality of the various supporting elements, or even how the task might depend upon development. Infant development studies (Smitsman, 1997) show that grasping is well established by 12 months and exists adequately even by 6 months. Thus we need to consider even earlier ages if we are to model the development of grasping. 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 guided. However, (Clifton et al., 1993) have performed infant reaching experiments *in the*

dark and shown that infants aged between 12 and 15 weeks are able to use proprioception alone, without vision, in successful reaching tasks. A question for investigation is thus the role of proprioception just before grasping behaviour has developed.

Consequently, if we follow our design principles, we should examine the earliest relevant stages first and this means investigating proprioceptive sensing in the limbs, *before* vision. We can also simplify as far as possible: it is sufficient to use only one of the arms and we employ a very simple fingertip touch sensor as the tactile sensing component. For the present experiments we fitted one arm with a simple probe or finger consisting of a 10mm diameter rod containing a small proximity sensor.

3. The Experimental Framework

We have now focussed our problem down into a first investigation into proprioception and the target system consists of a motor driven arm, sensory feedback on arm position, a contact sensor, and a simple spatial environment. This apparently limited structure provides the source of several important questions to be considered: proprioception encoding — how can space be sensed; motor control — how can new actions develop; coordination — intra-modal organisation; and constraint schedules — how should constraints be exploited?

3.1 Proprioception and tactile sensing

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. The question arises as to whether any particular form of neural encoding of limb position has more efficacy than others.

To illustrate the geometry of limb sensing, figure 2 shows a diagram of the configuration of the experimental robot arm. The arm consists of two limb sections, a "forearm" and "upper-arm" and the angles at the joints are given by θ_1 which is the angle between the upper-arm and the body baseline and θ_2 is the angle between the forearm and the the axis of the upper-arm.

This mechanical configuration is a 2 degree-of-freedom system and so we need 2 orthogonal variables to describe its state; let these be S_1 and S_2. The proprioception encoding question then concerns what could or should be the structure of the signals produced by these variables.

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

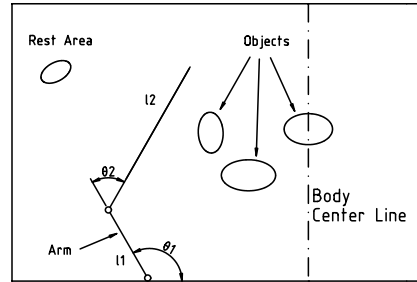$$S_1 = f(\theta_1) \qquad S_2 = f(\theta_2)$$

Figure 2: A plan view of the arm spatial configuration.

where f may exist as a linear or near linear scaling 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).

However, the joint angles have limited value for encoding the position of the arm, not least because the lengths of the limbs are changing significantly during infancy. 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. To locate end-points the proprioception signals could be as follows:

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

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

where l_1 and l_2 are the lengths of the upper-arm and forearm respectively, and 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 scheme would be to relate the arm end-points to the body centre-line. To obtain this **body-centred** encoding we simply use the shoulder encoding but shift the reference point from the shoulder to the centre of the body.

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 **Cartesian** encoding seems the most unlikely for a biological system, as it is the most abstract and has no obvious reference point to the

body. However we include this scheme due to its apparent importance in human spatial reasoning (Newcombe and Huttenlocher, 2000). We investigated all four systems as candidate encodings for proprioception signals.

3.2 An abstract motor model

In common with most animals, our robot arm has several distinct and independent degrees-of-freedom for movement. In theory, these could be realised as either revolute or linear jointed pairs but in biology joints are nearly always revolute. Each degree-of-freedom is usually powered by a muscle pair; known as extensors and flexors. We abstract this motor apparatus as a single motor parameter, M_i, which represents the overall drive strength applied to a given degree-of-freedom, i.

For the mechanics of actuation, we can assume that M_i represents the force applied by a given muscle pair and then a general equation of motion will include inertial, viscous friction and elastic components, thus: $M_i = k_1\ddot{\theta} + k_2\dot{\theta} + k_3\theta$, for joint angle θ_i, where k_1, k_2, k_3 are mass, viscosity and spring constants respectively. In fact, the viscous properties tend to be dominant, and so, to a first approximation, $M_i = k_2\dot{\theta}$ will determine the speeds of the limb segments being driven. For our 2-jointed arm we have two motor inputs, M_1 and M_2, real valued in the range $[-1, +1]$. By integrating the M_i over time during an action, each motor unit can output an "extent" signal, d_i, as a sense of motion.

For non-zero values of M_i the arm segments are started moving at constant speed and continue until either they reach their maximum extent or a sensory interrupt is raised. The ratio between the values of M_1 and M_2 determine the trajectory that the arm will take during an action. A small degree of noise is added to the motor system to create the effects of muscle tone. We assume that tone will increase with higher levels of attention and excitement, and also that motor noise reduces tone. Therefore we implement motor noise as decreasing in proportion to the excitation levels of target actions.

3.3 Mapping Sensory-Motor coordinations

Any model of sensory-motor learning will need some form of computational substrate upon which the relationships between sensory and motor events are experienced, learned and explored. We have developed such a computational substrate in the form of a mapping. 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 map structure: field size and inter-field spac-ing. These determine the degree of field overlap and the field density, per surface unit area. Every field in a map has an associated vector that can record local state information, this includes sensory stimulus values, excitation levels and repetition indicators. Cross-links between the fields of different maps record jointly experienced sensory-motor events and intra-modal and cross-modal coordinations.

We assume that uniform, un-assigned map structures are produced by prior growth processes but they are not pre-wired or pre-structured for any specific spatial system. Our system has to learn the correlations between its sensory and motor signals and the mapping structure is the mechanism that supports this. 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. Not all of the raw map may be used and the shape of the pattern of usage will depend upon the relationships and limits on the sensory-motor signals encountered. In this study we use uniform sheets of identical field sizes but we have also experimented with methods for growing fields of various sizes and locations on demand (in press).

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. It has been suggested that 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 with the width of view growing from 30 degrees at 2 weeks of age to 60 degrees by 10 weeks (Tronick, 1972).

Global excitation, which is the sum of the excitation levels over all the fields in a map, is an important parameter for controlling growth. In our experiments we used global excitation to trigger "constraint lifting". When global excitation becomes stable (i.e. flat plateau) then nearly all of the map has been ac-

cessed and another constraint can be removed. In this way, further map building may be started for another skill level.

3.5 Motivation and Attention

A motivational component is necessary to drive learning and there is evidence from infant studies that novelty is a strong motivational stimulus. In our limited system there are only two possible kinds of novel stimulus: a new field being initiated, or a change in a sensory value. Either of these will excite its relevant map field, increasingly so if it has not been experienced recently. Attention is then directed at the field with the highest level of excitation. Habituation mechanisms are used to reduce excitation with repetition and time, and so attention is attracted by novelty and decays with familiarity.

In order to initiate action we use a preset burst of activity as a reflex act. This rather ballistic approach to motor action is widely reported 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 but they do this by altering the timing and frequency of their actions, not the duration of the basic motor pattern (Thelen and Fisher, 1983). It seems that the neuronal burst duration is constant but the firing rate is modulated. This allows multiple muscles to be synchronised as they all have the same time-base while the amplitudes are varied to alter behaviour.

If the global excitation level is *very low* then the above reflex action is selected (with $M_1 = M_2 = +1$). If global excitation is *high* then the field with the highest level of excitation becomes the target for action. This target field and the field which corresponds to the current arm state are both accessed and their drive values, d_i, are retrieved. The differences between drive values are in proportion to the motor values that will produce a trajectory to the target field. The M_i values are thus computed using proportionality constants calibrated from previous movements.

However, there is a also probability of a purely random selection of motor values which increases in inverse proportion to the global excitation level. In other words, spontaneous actions can occur but only when global excitation is low, (i.e. when no stimulating activity has occurred for some considerable time).

A gross outline of the software system organisation is shown in figure 3 and further details are reported elsewhere (Lee and Meng, 2005).
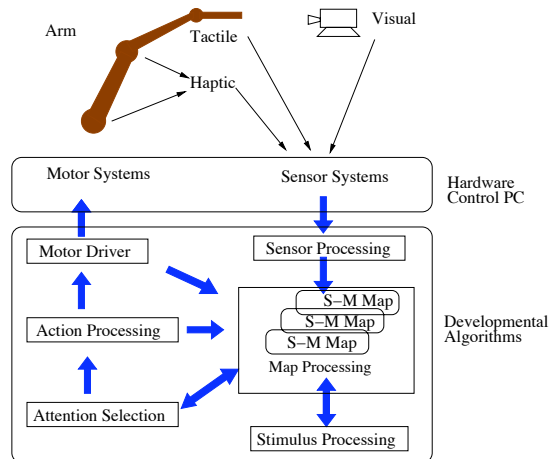

Figure 3: Software organisation.

4. Experiments and results

In this study of single modality proprioception we have examined four experimental variables. These are: the encoding scheme for kinaesthetic feedback, the degree of resolution of such feedback, the application of constraints, and the parameters that affect motivation and attention.

Regarding proprioception, we have four candidate encodings of the proprioceptive signals, S_1 and S_2 as described in Section 3.1 and each of these were tried in turn.

The resolution of proprioception is determined by the size of the fields used in the mappings. The effects of varying resolution were examined by creating three maps, each of different field size; small, medium and large, allowing all three to grow simultaneously. However, only one map can be used for attention and action selection, so by running the motors off each map in turn we observed the behaviour and effectiveness of the mapping parameters.

Regarding the effects of different field parameters, we used a number of different excitation and habituation regimes and various noise and weighting parameters. However, as attention is based on maximum excitation selection the only effect of altering the various habituation and decay parameters is to change the target stimulus for next action. This is seen as a preference for object B before A, rather than A before B, and, at this early stage of arm activity, such sequencing carries little significance.

For the application of constraints, the finger contact sensor was initially turned off. This meant a contact event did not interrupt movement and caused the system to sweep the hand from one extreme position to another, maybe pushing some objects out of the environment. This allowed the relation between the motor acts, M_i, the drive sense, d_i, and the extreme limits of sensed movement to be experienced and mapped. Then, when the contact sensor was

turned on, the stimulation of novel tactile contacts caused new spatial fields to be discovered. We also used the global state indicators to lift constraints in two other ways: actions are switched to finer resolution sensory maps when global familiarity gets very high, and the degree of spontaneous motor acts increases with very low global excitation.

4.1 Results

The first trials began with no contact sensing and no prior experience. Figure 4 illustrates the behaviour as traces of movements — for clarity these are displayed as directed lines between start and end fields in **motor drive space**. From this figure we see
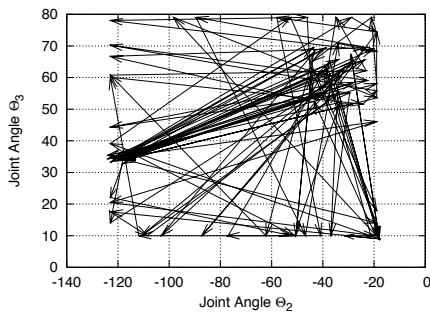

Figure 4: Arm movements with no contact sensing

that the arm first moved repeatedly to the body area (lower left), but as stimulation 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 for the sensory encoding in **Cartesian space** to show the locations in relation to the arm geometry.
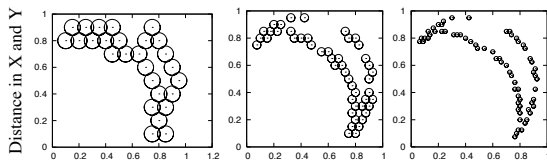

Figure 5: Fields generated during non-contact stage.

Eventually a plateau in field growth was reached and this was used as the trigger to lift the constraint on contact sensing. Figure 6 shows body moves now being interrupted by contact with an object on the hand path. Such contact events caused the exploration of the "internal" (non-boundary) fields. The effect of the first object contact is seen in figure 7 where some internal fields (and some more boundary fields) have been generated.

Each movement can be categorised according to the target field type: rest field, boundary field, or internal field. Figure 8 shows map growth in terms of these types and the characteristic plateau shape is evident.
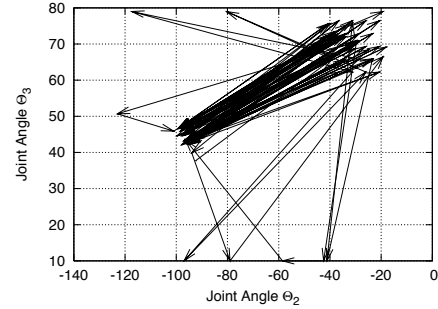

Figure 6: Arm movements with active contact sensing. An object (near the centre of the diagram) caused repeated contact action.
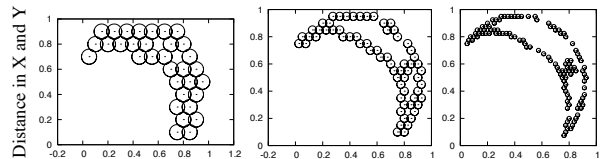

Figure 7: Fields generated after first object contact.

The behaviours observed from experiments form a progression: (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 "touching" cycles on detected objects. (5) directed touching of objects and sequences of objects. In the last case, a roughly cyclic behaviour pattern is produced, similar to eye scanpaths. All these behaviours, including motor babbling and the ballistic motor action, are widely reported in young infants (Piek and Carman, 1994).
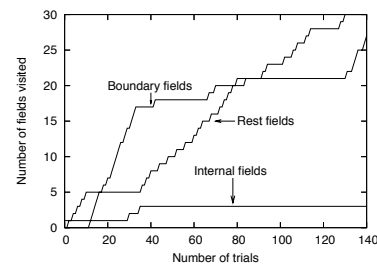

Figure 8: Rates of growth in maps. Only initial field visits are counted.

5. Observations

Having gained experience and data through experiments we can now revisit the key questions of proprioception encoding schemes, proprioceptive resolution, and the role of constraints.

We had expected that different encoding schemes would have a significant effect as they determine the shape of sensory-motor space as represented in the

mappings. However, there does not seem to be any clear advantage for any one of the four encoding schemes, at least for very early development. On reflection, all the schemes are continuous, smooth, nonlinear mappings and so they simply represent different distortions or warpings of the two-dimensional sheet. We recognise that when operating in the more kinematically restricted zones of the operating space there may be difficulties, but these are at the extremities where mobility is restricted anyway, and humans actually avoid these areas (Bernstein, 1967).

Without vision or any other feedback only the **joint angle** scheme seems likely to be effective immediately because the other three encodings require complex calculations and/or calibrations. 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 established. But, nevertheless, the other schemes could be developed through growth processes.

Although we found no compelling evidence for any one encoding scheme for the proprioceptive signals, we did make some notable observations. Firstly, when looking for the biological sources of proprioception we find not only joint angle sensors but also many muscle spindle receptors. These detect linear stretch in the muscles and it turns out that the anatomy of the muscles is arranged such that the spindle signals are even more appropriate for encoding space. If we assume that two limb segments are 20 units long and a muscle is affixed at 1 unit from the joint on one limb and 10 units from the joint on the other, then the length of the muscle will be: $\sqrt{101 + 20\cos\theta}$. When this, and the joint angle, θ, is plotted against the shoulder-hand distance, it is clear that the spindle gives a better linear approximation to the required distance measure than does the joint angle. In particular the spindle signal is most linear in the region where the joint angle is most difficult. This suggests that muscle stretch signals may play a more important role than joint angle receptors, at least in some configurations, and there are clear advantages if they coexist together, This has been reflected in research on the hind limbs of adult cats (Bosco et al., 2000) where it was 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.

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 both provide a **common** framework that can encompass both limbs. This is because they share a common reference point for the limbs and suggests the body centre-line has an im-
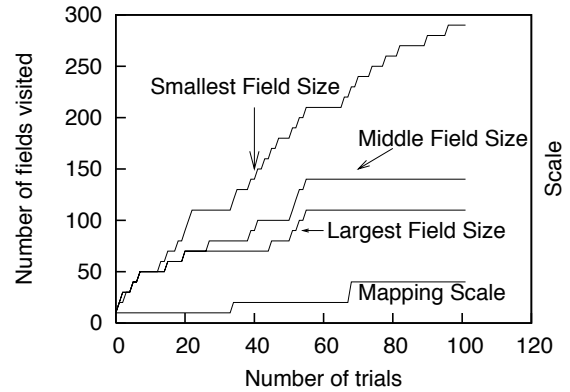

Figure 9: Transitions between three maps of different scale. Only initial field visits are counted. The "Mapping Scale" plot indicates the switching points.

portant role as a key fiducial point in egocentric space. Indeed, we notice that any head-mounted sensing system, such as the eyes, will also have a polar, body-centred reference frame which would also readily align with a body-centred proprioception system.

Regarding the resolution of the kinaesthetic sense we found 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 local 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 signalled by a global variable), the system transitioned to a finer scale map. Figure 9 shows this progressive transitioning over the 3 map sizes. This is supported by (Gomez, 2004) who has shown that starting with low resolution in sensors and motor systems, and then increasing resolution, leads to more effective learning. Also, the receptive field size of visual neurons in infants tends to decrease with age and development, and 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? As mentioned in section 2. visually guided reaching has long been believed to be the earliest accurate reaching behaviour to occur. A form of "hand looking" behaviour is almost certain 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 supports 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. Another model of infant grasping (Oztop et al., 2004) also suggests that visual guidance is not necessary for reaching but does not cover the growth of proprioception. Our model will be well placed to support the next stages of hand-looking and visual-guidance because 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 underemphazised the role of proprioception in infants' acquisition of prehension" (Clifton et al., 1993).

Acknowledgments

References

Ajemian, R., Bullock, D., and Grossberg, S. (2000). Kinematic coordinates in which motor cortical cells encode movement direction. *J. Neurophysiol*, 84:2191–2203.

Baraduc, P., Guigon, E., and Burnod, Y. (2001). Recoding arm position to learn visuomotor transformations. *Cerebral Cortex*, 11:906–917.

Bernstein, N. (1967). *The Coordination and Regulation of Movement*. Pergamon Press, Oxford.

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.

Braitenberg, V. (1984). *Vehicles: Experiments in Synthetic Psychology*. MIT Press, Cambridge, USA.

Bruner, J. (1990). *Acts of Meaning*. Harvard University Press, Cambridge, MA.

Clark, A. (1998). *Being There — Putting Brain, Body, and World Together Again*. MIT Press, Cambridge, MA.

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.

Gomez, G. (2004). Simulating development in a real robot. In *Proceedings of the 4th Int. Workshop on Epigenetic Robotics*, pages 1–24.

Hendriks-Jensen, H. (1996). *Catching Ourselves in the Act*. MIT Press, Cambridge, MA.

Lee, M. and Meng, Q. (2005). Psychologically inspired sensory-motor development in early robot learning. *International Journal of Advanced Robotic Systems*, 2(4):325 – 334.

Lungarella, M., Metta, G., Pfeifer, R., and Sandini, G. (2003). Developmental robotics: a survey. *Connection Science*, 15(4):151–190.

Mussa-Ivaldi, F. and Giszter, S. (1992). Vector field approximation: a computational paradigm for motor control and learning. *Biological Cybernetics*, 67:491–500.

Newcombe, N. and Huttenlocher, J. (2000). *Making Space*. MIT Press, Cambridge, MA.

Oztop, E., Bradley, N., and Arbib, M. (2004). Infant grasp learning: a computational model. *Experimental Brain Research*, 158:480 – 503.

Piek, J. and Carman, R. (1994). Developmental profiles of spontaneous movements in infants. *Early Human Development*, 39:109–126.

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.

Rutkowska, J. (1994). Scaling up sensorimotor systems: Constraints from human infancy. *Adaptive Behaviour*, 2:349–373.

Smitsman, A. W. (1997). The development of tool use. In Dent-Read, C. and Zukow-Goldring, P., (Eds.), *Evolving explanations of development*, pages 301–329. Am. Psy. Soc.

Thelen, E. (2000). Grounded in the world: Developmental origins of the embodied mind. *Infancy*, 1(1):3–28.

Thelen, E. and Fisher, D. (1983). The organization of spontaneous leg movements in newborn infants. *Motor Behaviour*, 15:353 – 377.

Thelen, E. and Whitmyer, V. (2005). Using Dynamic Field Theory to Conceptualize the Interface of Perception, Cognition, and Action. In Lockman, J. and Rieser, J., (Eds.), *Action as an Organizer of Learning and Development, Minnesota symposium on Child Psychology*, volume 33, pages 243–277.

Tronick, E. (1972). Stimulus control and the growth of the infant's effective visual field. *Perception and Psychophysics*, 11(5):373 – 376.

Westermann, G. and Mareschal, D. (2004). From parts to wholes: Mechanisms of development in infant visual object processing. *Infancy*, 5(2):131–151.