

## Aberystwyth University

### *Developmental Robotics from Developmental Psychology*

Lee, Mark; Meng, Qinggang; Chao, Fei

*Publication date:*  
2006

*Citation for published version (APA):*

Lee, M., Meng, Q., & Chao, F. (2006). *Developmental Robotics from Developmental Psychology*. 103-109. Paper presented at Towards Autonomous Robotic Systems, TAROS, Guildford, United Kingdom of Great Britain and Northern Ireland. <http://hdl.handle.net/2160/483>

#### **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](mailto:is@aber.ac.uk)

# Developmental Robotics from Developmental Psychology

M. H. Lee, Q. Meng and F. Chao

Department of Computer Science  
University of Wales, Aberystwyth, Wales, UK  
mhl@aber.ac.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 psychological knowledge of pre-grasping infants (up to 3 months of age) to examine the issues and factors that might produce similar mechanisms for use in robotic systems. The study includes discussion of results from robotic experiments on sensory-motor models and key issues are raised throughout.

## 1 Introduction

In the history of Artificial Intelligence, many avenues have been explored in the search for understanding of how truly autonomous learning robots might be created. In recent years these have included behaviour-based architectures (Brooks, 1991), hybrid approaches (Connell, 1992), evolutionary methods (Nolfi and Floreano, 2004), and probabilistic techniques (Thrun et al., 1998). Very recently rapid advances in brain science have led to increases in computational neuroscience and renewed interest in connectionist models. However, a missing element in all these approaches is the consideration of the developmental aspects of learning, that is, a treatment of learning in the context of an epigenetic framework that allows environmental and internal constraints to shape increasing competence and the gradual consolidation of control, coordination and skill (Prince et al., 2005). This situation is now changing, with new conference series such as (Epigenetics, 2005) and developmental robotics is now becoming established as a new strategy for robot learning, see (Lungarella et al., 2003) for a review.

In this paper we describe an approach to sensory-motor learning and coordination that draws from psychology rather than neuroscience. There have been many models of sensory-motor coordination (Lungarella et al., 2003) but most of these have been

based on a specific, behavioural task. We are interested in exploring mechanisms that can support not only the growth of behaviour but also the transitions that are observed as behaviour moves through distinct stages of competence. It is essential that this is examined logically if we are to understand the options and requirements for algorithms that could support developmental learning in machines. Following the methodology of (Thelen and Whitmyer, 2005) we adopt a “content-neutral” approach in which we strive for general rather than task-specific models, and avoid assumptions about internal belief states or internal causal knowledge. As space is limited, this paper is descriptive (full technical detail is available from the authors, and in forthcoming papers) in order to explain the approach and its implications.

## 2 Developmental Learning

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. We are interested in very early development, in particular, the control of the limbs and eyes during the first three months of life. The newborn human infant faces a formidable learning task and yet advances from undirected, uncoordinated, apparently random behaviour to eventual skilled control of motor and sensory systems that support goal-directed action and increasing levels of competence. This is the kind of scenario that will face future robots and we need to understand how some of the infant’s learning behaviour might be reproduced.

A major source of inspiration for the developmental approach comes from the prominent psychologist Jean Piaget (Piaget, 1973) and we recognise Piaget’s emphasis on the importance of sensory-motor interaction, staged competence learning and the sequential lifting of constraints (or scaffolding) (Piaget, 1973). Others, such as Jerome Bruner, have further studied the plasticity seen in infant studies and developed Piaget’s ideas further, suggesting mechanisms that could explain the relation of symbols to motor acts, especially concerning the

manipulation of objects and interpretation of observations (Bruner, 1990).

The first year of life is a period of enormous growth and various milestones in development are recognised. Considering just motor behaviour, newborn infants have little control over their limbs and produce uncoordinated and often ballistic actions, but over the first 12 months, control of the head and eyes is gained first, then follows visually guided reaching and grasping, and then locomotion is mastered. 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). We believe it is necessary to begin investigations at the earliest levels of development because early experiences and structures are likely to determine the path and form of 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).

### 3 The Sensory-Motor Coordination Problem

Even before any cross-modal spatial integration can begin it is necessary to first discover the structure of the local spaces within each modality. By virtue of their differing physical structure and constraints, each modality will have its own coding of space. Thus, when the 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 hand; for example, locations in end-effector space are encodings of signals that correspond to the hand being at a certain location.

During the first months of life the neonate may seem to show no purpose or pattern in motor acts, but actually the infant displays very considerable adaptation: from spontaneous, 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). Various stages in behaviour can be discerned, during which the local egocentric limb space becomes assimilated into the infant’s awareness and forms a substrate for future cross-modal skilled behaviours. This essential correlation between proprioceptive space and motor space seems to be a foundation stone for development, and occurs at many levels (Pfeifer and Scheier, 1997). Sensory-motor coordination in the limbs appears to precede visual development (it may begin in the womb). The very early stage of infant growth does not rely on vision (Piek and Carman, 1994), and even when it can continue concurrently with visual development, in the first few months, the eye is too functionally restricted to correlate closely with other modalities

(Westermann and Mareschal, 2004). We believe it is important to start with the very earliest stage of coordination of the limbs, and so in the experiments reported here we do not involve an eye system.

#### 3.1 An abstract motor model

In general, the skeletal system of an animal will allow movement through many distinct and independent degrees-of-freedom. In theory, these could be realised through 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 viscous friction and elastic components, thus:  $M_i = k_1\ddot{\theta} + k_2\dot{\theta} + k_3\theta$ , for joint angle  $\theta_i$ . 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. We believe this is a reasonable abstraction for many practical purposes.

#### 3.2 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 1 shows a diagram of the configuration of one of our experimental robot arms. This consists of two limb sections, a “forearm” and “upper-arm” and the angles at the joints are given by  $\theta_1$  which is the angle between the upper-arm and the body baseline and  $\theta_2$  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 these state 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 near 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




Figure 1: A plan view of the arm spatial configuration.

by the afferent fibres from mechanoreceptors embedded in the joint capsules (Bosco et al., 2000).

However, when the lengths of the limbs are changing due to growth then joint angles become insufficient to describe the position of the limb end-points and other, more complex, possibilities can be considered. To locate end-points the proprioception signals could be as follows:

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

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

where  $l_1$  and  $l_2$  are the lengths of the upper-arm and forearm respectively, 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, even more attractive 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 reference point to the body. However we include this scheme due to its 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.

## 4 A Mapping Model for Sensory-Motor 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 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 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, excitation levels and motor values. The field values in a map decay with time and can be viewed as a form of short term memory.

We assume that basic uniform 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 consistent field sizes but we have also experimented with methods for growing fields of various sizes and locations on demand.

## 5 Reflex Action and Constraint Lifting

A motivational component is necessary to drive learning and there is evidence from infant studies that novelty is a strong motivational stimulus. We define any new sensory value that has not been experienced recently as a novel stimulus that excites relevant map fields. Attention is then directed at 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.

In order to initiate action we use a preset burst of activity as a reflex act. By driving the motors ‘full on’ brings the hand from the lateral rest position to the body centre-line. 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.

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 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).

Our experiments are aimed at exploring the value of constraints in facilitating learning. In this approach, ‘constraint lifting’ is the key to progression towards increasing competence. We use global excitation parameters to trigger constraint lifting. For example, the sum of the stimulated fields will reach a plateau when all of the map has been accessed and processed. In this way, further map building may begin for another skill level, thus simulating a form of Piagetian learning. In this study we use global state indicators to lift constraints in two ways: finer resolution sensory maps are used when

global familiarity is high, and the degree of spontaneous motor acts increases with very low global excitation.

## 6 Experimental Results

An experimental laboratory robot system was used to implement the above concepts. Figure 2 shows the general system in which one arm is configured as in figure 1. Due to lack of space, the software system, illus-



Figure 2: Part of the experimental robot system.

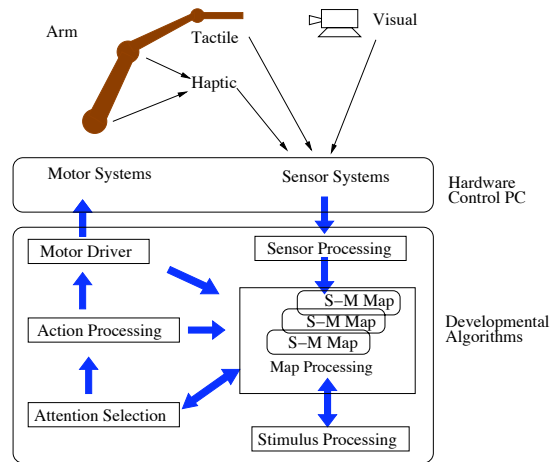


Figure 3: Software organisation.

trated in outline in figure 3, is described in more detail elsewhere (Lee and Meng, 2005) and in press. For more details of the excitation and habituation models used see (Meng and Lee, 2005).

Following our approach, which is described in more detail in (Lee et al., 2006), the experiments were designed from a logical examination of the possible parameters that may be varied.

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 fields in motor space. From this figure we see that the arm moved re-




Figure 4: Arm movements with no contact sensing.

peatedly between rest and body areas first, 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 in Cartesian space to show the locations in relation to the arm geometry. Eventually a plateau in

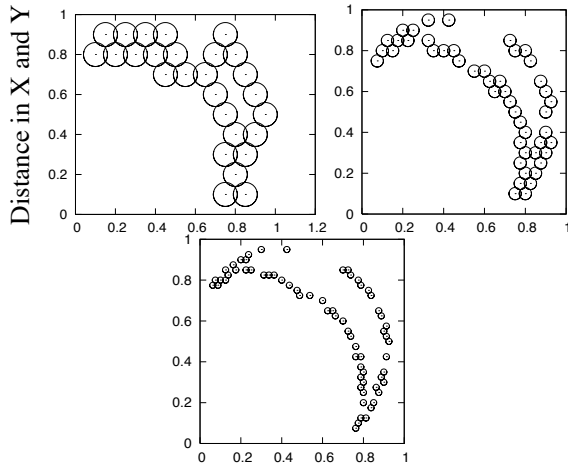


Figure 5: Fields generated during non-contact stage.

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 created many internal fields.

Each movement can be categorised according to the target field type: rest field, boundary field, or internal (contact) field. Figure 8 shows map growth in terms of these types and the characteristic plateau shape is evident. It is notable that the pattern of fields visited at the beginning is very different from that produced after most fields have been visited.

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 cycles of contact




Figure 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.

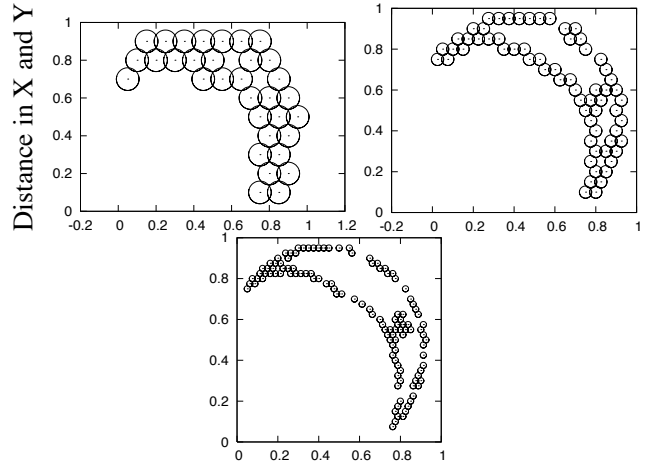


Figure 7: Fields generated after first object contact.

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. All these behaviours, including motor babbling and the rather ballistic approach to 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.

## 7 Discussion

Regarding proprioception, initially there does not seem to be any clear advantage for any one of the four encoding schemes. All the schemes are continuous, smooth non-linear mappings and so they simply represent different distortions or warpings of the two-dimensional sheet. Figure 9 shows the way in which the mappings grow




Figure 9: Map correlation. Each field in the sensory map (upper) is connected to a field in the motor map (lower).

— the upper layer is the map of kinaesthetic feedback in Cartesian encoding and the lower layer is the motor drive map. This shows just boundary fields (the internal areas become filled through object contact experience) but the warped correlation between the maps is

clearly noticeable. 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 and humans actually avoid these areas (Bernstein, 1967).

Although we found no compelling experimental 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 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. However, it turns out that the anatomy of the muscles is arranged such that the spindle signals are even more appropriate for the spatial encodings. 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 is plotted against the shoulder-hand distance as the joint rotates through 180 degrees it is clear that the spindle gives a more 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 advantages if they coexist together, thus reinforcing the findings of Boscoe and Popple (Bosco et al., 2000).

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 includes both limbs. This is because they have one common reference point for the limbs, not two, and suggests the body-centre has an important 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 very easily align with a body-centred proprioception system.

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 10 shows this




Figure 10: Transitions between three maps of different scale. Only initial field visits are counted. The “Mapping Scale” plot indicates the switching points.

progressive transitioning over the 3 map sizes. 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 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).

## 8 Relation with other work

The majority of previous models of sensory-motor coordination have been based on connectionist archi-

tectures (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 also tend to use very large numbers of neuronal elements. As a review commented: “their behavioural capacity is usually limited” (Kalaska, 1995). While “motor babbling” is seen in the behavioural output of many systems, very few follow the psychological literature on development and even less deal with transitions between qualitatively distinct patterns of behaviour. Models of infant grasping have been produced (Oztop et al., 2004) and although they also suggest that visual guidance is not necessary for reaching, they do not cover the growth of proprioception.

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 tries to cross-correlate all possible events and is 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). Most of these models are based on the concept of a sensory-motor *schema* drawn from Piaget’s conception of schemas in human activity (Piaget, 1973). The linked fields in our maps can be viewed as elemental schemas that relate local sensory and motor events.

## Acknowledgments

We thank R.E. Gunstone for developing parts of the software implementation. We are grateful for the support of EPSRC through grant GR/R69679/01 and for laboratory facilities provided by the Science Research Investment Fund.

## References

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



- point based or joint-level based? *J. Neurophysiol.*, 83(5):2931–45.
- Brooks, R. (1991). Intelligence without representation. *Artificial Intelligence*, 22.
- Bruner, J. (1990). *Acts of Meaning*. Harvard University 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.
- Connell, J. (1992). SSS: A hybrid architecture applied to robot navigation. In *IEEE Conference on Robotics and Automation (ICRA-92)*, pages 2719–2724.
- Drescher, G. (1991). *Made up minds: a constructivist approach to artificial intelligence*. MIT Press, Cambridge, MA.
- Epigenetics (2001-2005). *Proceedings of the Int. Workshops on Epigenetic Robotics*.
- 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, D. (1982). *Understanding Motor Development in Children*. John Wiley, NY.
- 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.
- 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.
- 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.
- Lee, M. H., Meng, Q., and Chao, F. (2006). A content-neutral approach for sensory-motor learning in developmental robotics. In *Proceedings of the 6th Int. Workshop on Epigenetic Robotics*.
- Lungarella, M., Metta, G., Pfeifer, R., and Sandini, G. (2003). Developmental robotics: a survey. *Connection Science*, 15(4):151–190.
- Meng, Q. and Lee, M. (2005). Novelty and habituation: the driving forces in early stage learning. In *Neural learning for intelligent robotics*, pages 315–332. LNCS.
- Newcombe, N. and Huttenlocher, J. (2000). *Making Space*. MIT Press, Cambridge, MA.
- Nolfi, S. and Floreano, D. (2004). *Evolutionary Robotics*. MIT Press, Cambridge, MA.
- Oztop, E., Bradley, N., and Arbib, M. (2004). Infant grasp learning: a computational model. *Experimental Brain Research*, 158:480 – 503.
- Pfeifer, R. and Scheier, C. (1997). Sensory-motor coordination: the metaphor and beyond. *Robotics and Autonomous Systems*, 20(2):157–178.
- Piaget, J. (1973). *The Child’s Conception of the World*. Paladin, London.
- Piek, J. and Carman, R. (1994). Developmental profiles of spontaneous movements in infants. *Early Human Development*, 39:109–126.
- Pouget, A. and Snyder, L. (2000). Computational approaches to sensorimotor transformations. *Nature neuroscience*, 3:1192–1198.
- Prince, C., Helder, N., and Hollich, G. (2005). On-going 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.
- 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.
- Thrun, S., Burgard, W., and Fox, D. (1998). A probabilistic approach to concurrent mapping and localization for mobile robots. *Autonomous Robots*, 5(3-4):253 – 271.
- 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.