Aberystwyth University

Staged Competence Learning in Developmental Robotics

Lee, Mark; Meng, Qinggang; Chao, Fei

Published in:
Adaptive Behavior

DOI:
[10.1177/1059712307082085](https://doi.org/10.1177/1059712307082085)

Publication date:
2007

Citation for published version (APA):

Lee, M., Meng, Q., & Chao, F. (2007). Staged Competence Learning in Developmental Robotics. *Adaptive Behavior*, 15(3), 241-255. <https://doi.org/10.1177/1059712307082085>

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

Staged Competence Learning in Developmental Robotics

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

*Department of Computer Science,
University of Wales, Aberystwyth, Ceredigion, SY23 3DB, Wales, UK

**Department of Computer Science,
Loughborough University, Leicestershire, UK

Contact: mhl@aber.ac.uk

Abstract

Developmental psychology has long recognised the presence of stages in human cognitive development, although the underlying causes and processes are still an open question and subject to much debate. This paper draws inspiration from psychology and describes an approach towards developmental growth for robotics that utilises natural constraints in a general learning mechanism. The method, summarised as Lift-Constraint, Act, Saturate (LCAS), is applicable to all levels of control and behaviour, and can be implemented in any robotic configuration. An implementation based on sensory-motor learning in early infancy is described and the results from experiments are presented and discussed.

Keywords: Developmental robotics, Sensory-motor control, Constraint-based learning.

1. Introduction

Truly autonomous robots must be able to operate and survive unaided over reasonably long operational periods. But this remains a difficult research challenge despite the ample evidence of autonomy in biology. One key reason is that quite sophisticated adaptation or learning abilities are essential because in real, unstructured environments novel experiences are unavoidable and must be handled appropriately. This means that the system must not only adapt in accordance with current experience but must also be capable of adapting its learning processes themselves. Thus, new learning processes must emerge as conditions change and new demands are made.

The way this problem has been solved in the more cognitively advanced biological systems is through processes of structured growth generally known as “development”. It is surprising that only recently has any real attention been given to the concept of development as an important factor for autonomous systems. Despite the vast body of research knowledge on learning and adaptation, in both animals and machines, built up over the last five decades, very little work has considered implementing developmental processes in artificial systems. This is remarkable because psychologists have studied development for much longer and the idea is not new or original. Indeed, Alan Turing actually suggested this in 1950:

“Instead of trying to produce a programme to simulate the adult mind, why not rather try to produce one which simulates the child’s? If this were then subjected to an appropriate course of education one would obtain the adult brain. ... We have thus divided our problem into two parts. The child programme and the education process. These two remain very closely connected.” (Turing, 1950)

Turing did not give a starting age for the child program, (he said “Opinions may vary as to the complexity which is suitable in the child machine”) and the “course of education” can be taken in its broadest sense to include experience in general, but it is clear that he is talking about cognitive development:

“In the process of trying to imitate an adult human mind we are bound to think a good deal about the process which has brought it to the state that it is in.” (Turing, 1950)

The continuing growth of research in brain science has stimulated much activity both in biologically-targeted

computer modelling and in adapting biological concepts for the design and development of novel computing technologies. However, developmental aspects, in particular cognitive growth, have only recently been given much attention in robotics. The topic of developmental robotics is now becoming established as a new research area for robot adaptation, control and learning (Lungarella, Metta, Pfeifer, & Sandini, 2003). This approach emphasises the role of environmental and internal factors in shaping adaptation and behaviour, and posits a developmental framework that allows the gradual consolidation of control, coordination and competence (Prince, Helder, & Hollich, 2005).

The challenge from this viewpoint is in finding effective algorithms for processes that support the development of learning and adaptation. There exists a large gap between our psychological theories of development and our ability to implement working developmental algorithms in autonomous agents. This paper describes an approach towards closing that gap and illustrates with an experimental implementation. The next section introduces the concept of staged growth and then Section 3 shows how this depends upon constraints. Section 4 outlines our method and then follows a description of an experimental robot system. Section 6 then shows how the method is used to organise a schedule of constraints, and Section 7 presents the experimental results. The final sections discuss the implications, related work, and conclusions.

2. Developmental Stages

Developmental psychology concerns the study of human cognitive growth and the production of theories that could explain such growth. Experimental findings inform the theories and the data is inevitably based on patterns of behaviour and changes in behaviour over time. This focus on behaviour means that, unlike neuroscience for example, there is no structural information that might be used to drive inferences about possible internal mechanisms of adaptation that could account for the manifest behaviour. Consequently, there have been many contrasting theories of development in psychology.

A key characteristic of animal development is the centrality of behavioural sequences: no matter how individuals vary, all infants pass through sequences of development where some competencies always precede others. This is seen most strongly in early infancy as one pattern of behaviour merges into another. These regularities are the basis of the concept of behavioural stages — periods of growth and consolidation — followed by transitions — phases where new behaviour patterns emerge. Perhaps the most influential theories of staged growth have been those of Jean Piaget who emphasised the importance of sensory-motor interaction, staged competence learning and a constructivist approach (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 (Kalnins & Bruner, 1973; Bruner, 1990). Very briefly, Piaget defined four periods in an individual's life: beginning with the Sensorimotor period (up to 2 years) during which infants are not fully capable of symbolic representation; then the Preoperational period (2 to 6 years) is characterised by egocentric behaviour; then follows the Concrete Operational period (6 to 12 years) in which abilities in classification and linear ordering are seen; and finally, the Formal Operation period (from 12 years onwards) displays capabilities in formal, deductive, and logical reasoning.

All periods and stages have vague boundaries and overlap and merge; they are not fixed but vary greatly with individuals. There is also a great deal of debate about the origin and drivers for developmental change. At one extreme, Nativism argues that the machinery for development is genetically determined and the processes of growth are largely preprogrammed, with any acquired cognitive competence being either ignored or refuted. In stark opposition, Empiricism takes the view that experience is a major factor in shaping the course of development and that structures acquired to support development are shaped by experience. This debate, also known as the Constructivism versus Evolutionism argument, still continues after many decades and has various implications for psychology theory. Nevertheless, the existence of stages in development is not controversial and their role in the growth of cognition appears to be very significant.

We believe that research into developmental algorithms for robotics should begin with and be firmly grounded in the sensorimotor period. This is for several reasons: (1) it is logical and methodologically sound to begin at the earliest stages because early experiences and structures are highly likely to determine the path and form of subsequent growth in ways that may be crucial; (2) according to Piaget, the sensorimotor period consists of six stages that include concepts such as motor effects, object permanence, causality, imitation, and play — these are all issues of much relevance to robotics; (3) sensory-motor adaptation is vital for autonomous robots as motor behaviour has to be mastered and changed to suit changing circumstances in the environment; (4) it seems likely that sensory-motor coordination is a significant general principle of cognition (Pfeifer & Scheier, 1997).

For these reasons we focus on the earliest level of sensory-motor development: the emerging control of the limbs and eyes during the first three months of life. To the casual observer the newborn human infant may seem

helpless and slow to change but, in fact, this is a period of the most rapid and profound growth and adaptation. 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). We believe there is much to learn for robotics from this scenario.

3. The Importance of Constraints

All processes of adaptation and learning must have some forms of underlying bias or a priori assumptions. This is because choices have to be made in representations, learning approach, possible actions, etc, even before learning begins. These biases are treated in developmental psychology in terms of constraints and there are many theories as to the origins, role and effects of such constraints on development. See (Keil, 1990) for a review of constraints under different theories.

Any constraint on sensing, action or cognition 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). When a high level of competence at some task has been reached then a new level of task or difficulty may be exposed by the lifting of a constraint (Rutkowska, 1994). The next stage then discovers the properties of the newly scoped task and learns further competence by building on the accumulated experience of the levels before.

Various examples of internal sensory and motor constraints are seen in the newborn, for example the neonate has a very restricted visual system, with a kind of tunnel vision (Hainline, 1998) where the width of view grows from 30 degrees at 2 weeks of age to 60 degrees at 10 weeks (Tronick, 1972). Although this may seem restricted, these initial constraints on focus and visual range are “tuned” to just that region of space where the mother has the maximum chance of being seen by the newborn. When “mother detection” has been established then the constraint can be lifted and attention allowed to find other visual stimuli.

Many forms of constraint have been observed or postulated (Hendriks-Jensen, 1996) and we can consider a range of different types:

Anatomical/Hardware These are the physical limitations imposed by the morphology of an embodied system. These include kinematic restrictions from structural (e.g. skeletal) joints and spatial configurations. Also mechanical constraints may be relevant, such as motor limitations preventing freedom of movement.

Sensory-Motor All sensors have their own limitations; usually these are specified in terms of accuracy, res-

olution, response rates, and bandwidth. Motor systems also have similar characteristics with additional features for dynamic performance.

Cognitive/Computational Constraints on cognition take many forms, not just relating to speed but also relating to information content and structure. Many constraints that affect artificial neural systems are now known, and it is likely that some of these exist in human cognition. Sensory constraints can also affect or limit the input to cognitive processes.

Maturational These are the most difficult to enumerate but it is certain that internal biological growth processes influence and maybe facilitate cognitive growth. Both neural and endocrinal support systems will have effects as they mature, for example neurogenesis is still underway for much of early infancy.

External/Environmental External constraints are those that restrict behaviour or sensory input in some way but originate from the environment not from the individual or agent. This is a very powerful source of constraint as they can be applied at any time and are not related to the individual’s stage of growth. Examples of this are seen in the use of carefully structured environments by another agent to assist learning and adaptation. Thus, in parental care, education and many other social situations, interactions generate patterns in the environment in order to direct attention or action towards a goal. This is known as scaffolding.

Much of the developmental literature concerns the role of constraints in higher order cognitive tasks such as number, language and reasoning (Campbell & Bickhard, 1992). These internal, cognitive constraints deal with issues like representation and could be termed “soft” constraints. We are interested in the most basic processes of sensory-motor adaptation and so are concerned with the “hard” constraints that emerge from the actual physical properties and features of the system. These strongly influence the construction of the adaptive processes and could be considered as Type 4 constraints (Campbell & Bickhard, 1992).

4. The LCAS approach

To properly investigate the options and requirements for developmental algorithms it is necessary to adopt a logical approach. Following the methodology of (Thelen & Whitmyer, 2005) we follow 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, (Lee, Meng, & Chao, 2006). We emphasise explicit, abstract models and try to avoid preselected internal representations.

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 seems that the trigger for transitions must be related to internal global states, not local events, because transitions generally only occur when there is no local activity.

For example, if a novel object raises some local excitation levels, then successive new stimuli will cause increasing global excitation (global excitation being the sum of all local excitation) but global excitation will only reach a stable plateau when no novel events have been seen for some time. Thus, high competence at a level is equivalent to all incoming experience matching that expected, with no novel changes or unexpected events.

Thus, global states such as global excitation can act as indicators that can detect qualitative aspects of behaviour such as when growth changes have effectively ceased or when a map has become saturated. They can then signal the need to enter a new level of learning by lifting a constraint or accessing a new sensory input. In this way, further exploration may begin for another skill level, thus approximating a form of Piagetian learning.

Our approach then consists of implementing the cycle; Lift-Constraint, Act, Saturate, at a suitable level of behaviour. First, the possible or available constraints must be identified and a schedule or ordering for their removal decided. Next a range of primitive actions must be determined together with their sensory associations. Also any sensory-motor learning or adaptation mechanism is incorporated at this stage. Finally a set of global measures need to be established to monitor internal activity. When this is implemented the initial behaviour may seem very primitive, but this is because all or nearly all constraints have been applied and there is little room for complex activity. Then the Act stage begins and through varying patterns of action the scope for experience is thoroughly explored and all new experiences are learned and consolidated. Eventually there are no new experiences possible, or they are extremely rare, and this level becomes saturated. The global indicators then reach a critical level and the next constraint in the schedule is lifted and the cycle begins again.

5. An Experimental System

We use a hand/eye hardware system as a test-bed for exploring and assessing different developmental algorithms. This laboratory 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 color camera, is mounted above and looks down on the work area. Figure 1 shows a view of one of the robot arms.

Figure 1 here.

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 . There is a lateral reference position, known as the rest area, and another area, known as the body area, is reached by driving both motors full on from the rest position.

Figure 2 here.

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. The learning should be fast, incremental and constructive, and we have designed an appropriate computational framework for sensory-motor coordination.

We use a mapping scheme as a substrate for capturing and processing sensory-motor experience. 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**. Each field represents a region of a sensory or motor space that is perceptually equivalent, that is, all points within a field are perceptually indistinguishable. The fields are circular, regularly spaced, and overlapping. The degree of field overlap and the field density, per surface unit area, are defined by two parameters: field size and inter-field spacing. Every field in a map has a set of associated variables that can record state information, these include sensory stimulus values and excitation levels.

Each map layer deals with a single channel of either sensory 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 (showing the limits of the working envelope). 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 distortion or warping is clearly visible.

Figure 3 here.

There are various possible ways of creating fields and we have experimented with methods for growing fields of various sizes and at various locations on demand (Meng & 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 — they have no *intrinsic* relation with any external space. Thus, a map starts as an empty sheet, and the fields gradually become populated with sensory or motor data for experiential events. The shape of the pattern of usage will depend upon the relationships between the signals encountered.

In order to motivate action we use a simple novelty detector. Any new field or any field that has different stimulus values from the current sensory input values are considered as unexpected events (novel) and the field excitation variables are increased. Habituation functions are applied to all excitation levels and so the field values in a map decay with both repetition and time. Thus, attention is attracted by novelty and decays with familiarity.

From the local variables in each field we can compute some global state values for each map. Global excitation, G_e , is a measure of total excitation and is the sum of the excitation levels of all the fields whose excitation levels are above a nominal lower threshold. Global conversancy, G_f , is a measure of the state of development of a map. Each field has a frequency value and initially all these are set to zero. Then each time a field is accessed its frequency value is incremented. Consequently, a summation of all the frequency values will indicate how much of the map has been created and how well used (or familiar) it has become. G_f is designed to give a decreasing value as the fields become less novel and increasingly explored. Global excitation can be seen as an indication of the intensity of current activity and global conversancy is a measure of the novelty or newness of the fields being experienced. For further details of the implementation and software organisation see (Lee & Meng, 2005).

6. Constraint Analysis

We can now consider the experimental system from a constraints perspective and decide how the experiments could take advantage of a constraint lifting schedule in order to enhance learning.

The system consists of a physical anatomy (or embodiment), an environment and a set of very basic sensory-motor functions. Following section 3., we can identify all the constraints available in our particular system:

Anatomical structure — It was a deliberate design decision to arrange for only 2 degrees of freedom in each arm. With more degrees of freedom the issue of redundancy becomes a problem for motor control (Bernstein, 1967) but this has been covered elsewhere (Sporns & Edelman, 1993). This constraint was fixed throughout.

Kinaesthetic or visual priority — Rather than let all behavioural modalities develop simultaneously there may be advantages in imposing some order. Sensory-motor coordination in the limbs appears to precede any significant visual development (it seems to begin in the womb). This very early stage of infant growth does not rely on vision (Piek & 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 very closely with other modalities (Westermann & Mareschal, 2004). We adopted the constraint that vision should not be active (i.e. not subject to adaptation or learning) while the limb kinaesthetic sense was developing. This was fixed throughout.

Motor activity — The basic arm action is parameterised by two motor values, M_1 and M_2 , one for each limb segment, and the ratio between the values determines the trajectory that the arm will take. The limb segments move at constant speed and continue until either they reach their maximum extent or a sensory interrupt is raised. This ballistic approach to motor action is readily seen in limb action in infants, and by three months they also have considerable proprioceptive control of their legs and without any visual input can learn different action patterns (Angulo-Kinzler, Ulrich, & Thelen, 2002).

We can introduce various constraint possibilities into this basic act. If the global excitation level is not low then the field with the highest excitation level will be selected as the target for action. This is a simple action selection function. If global excitation is *very low* then there is no target for action and a reflex action is performed that moves the hand to the body area. However, there is a small probability of a purely random selection of motor values to create a spontaneous action. This is arranged so that the likelihood of spontaneous motor acts increases with very low global excitation.

Sensory resolution — The resolution of the proprioceptive system is a possible constraint variable. We

use different sizes of fields in the maps to produce different resolutions. Coarse resolution must come before fine so we can arrange that constraints on finer resolution are lifted after the coarser cases have been explored first. The global familiarity indicator can be used to indicate when a map has been fully explored and is ready for the next level of detail.

Sensory sequencing — There is only one choice in our current experiment; should the proprioceptive system be learned before tactile sensing, or vice versa? It is clear that there is no point in using tactile before proprioception as contact without spatial localisation is meaningless. On the other hand, it might be advantageous to discover spatial structure before object perception and so a constraint on tactile sensing was adopted for experimentation.

Environmental constraints — The scope for constraints here simply consists of the presence or absence of objects. The shape or other properties of objects is not relevant as the tactile sensor resolution is restricted to simple contact events.

7. Experiments and results

The experiments were designed to probe the role of the constraints described above. This includes active and inactive contact sensing, the use of objects, and sensory resolution. In addition to the constraints identified above, there are several system parameters that could affect performance and are available for experimentation, these include: different proprioception spatial encoding schemes, different motor action patterns, and variations in excitation/habituation parameters. The first two are reported in (Lee & Meng, 2005) and the latter in (Meng & Lee, 2005), but it is important to note that none of these have any significant effects on the results reported here. Their main effect is to alter the order or sequencing of local behaviour but the coordination structures built over the experiments and the general behavioural patterns were very similar.

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 sensory-motor data were used to build each map separately and simultaneously. However, attention and action selection can only choose one field, so actions were driven from each map in turn in separate experiments.

7.1 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.

Figure 4 here.

From 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.

Figure 5 and 6 here.

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 path. Such contact events create internal (non-boundary) fields, as seen in figure 10., and generate much repeated action.

Figure 7 here.

Figure 8 illustrates field growth rates 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 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.

Figure 8 here.

From the experiments we observe a progression of qualitatively distinct behavioural patterns:

“Blind groping” actions mainly directed towards the body area,

Extended groping directed towards the bounding limits of the agent's egocentric space,

Unaware contact seen as pushing or ejecting objects out of the local environment,

Contact sensitivity where limb movements are interrupted by tactile sensing events,

Repeated cycles of contact observed as repeated “touching” behaviours directed at detected objects,

Directed touching of objects and sequences of object touching cycles.

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.

Figure 9 here.

All these behaviours, including motor babbling and the rather ballistic motor actions, are widely reported in young infants (Piek & Carman, 1994).

8. Discussion

Our choice of constraining visual development until after a kinaesthetic sense has been established could be controversial but the results show that this is not an unreasonable developmental sequence. Much of the psychological literature tends to assume that vision is the dominant sense and 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, Muir, Ashmead, & Clarkson, 1993) have performed infant reaching experiments in the dark and shown that infants of around 15 weeks are able to use proprioception alone, without vision, in successful reaching tasks. A form of “hand looking” behaviour can be expected 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* occur prior to visual development, can be used to guide action, and does not necessarily depend upon visual confirmation. A well developed kinaesthetic sense could be a great advantage in supporting visual-guidance and visual coordination by providing a ready mapping of the local operating 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).

The constraint on motor acts has been rather weaker than the others used here. It consisted of allowing spontaneous actions to occur occasionally when global excitation became very low. This can be seen as creating a disturbance to generate some potentially new experiences whenever there has been a prolonged lack of interesting activity. This device could be viewed equally as a constraint derived from a developmental process or as part of the internal mechanism of the motor system.

Regarding the effects of sensory resolution on proprioception, 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. This effect was also reported by (Gomez, 2004). Figure 10 shows the transitions across the maps, from coarse to fine, as plateaus in their activity were detected. The transitions occurred at trials 34, to medium fields, and at 68 for the smallest field maps. 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 & Mareschal, 2004).

Figure 10 here.

The integration of proprioception and tactile senses can produce a powerful haptic system but it is an open question as to which part should develop first. Our speculation that tactile sensing could be delayed until after significant kinaesthetic growth in the same modality is seen to be supported by the results. At least, it is a viable strategy to reduce the complexity of the learning input by discovering some of the structure of local space before tactile sensing data is processed. Of course, a very complex tactile system such as the hand with many types of receptors sensing heat, vibration, pain and touch may well need a period of familiarisation to establish the various functions but this is distinct from object detection, and could take place in parallel with other activities. Another possibility is that both components of the haptic system develop together; this happens in our experiments but only if the tactile events are supported by relevant spatial learning — the tactile system cannot lead because it relies on a spatial frame as a context for its experiences.

Regarding environmental constraints, we could only use the idea of scaffolding to the extent that we could place objects in areas that were under-explored and thus direct attention to gain developmental experience in those areas. This was only possible after the tactile constraint had been lifted; before then objects would be ignored and possibly ejected. In later work we have examined the effects of known objects being removed, and this leads on to object permanence and the detection of moving objects and external agency.

9. Relation with other work

While there is now growing research activity in the area of developmental robotics, most such work deals with specific topics such as motivation, active vision, self-awareness, interaction, and modelling issues. Much of

this research has relevance for our approach, as seen in the citations given, particularly those that shed light on possible mechanisms and algorithms. But there is still a lack of research that takes account of the large body of experimental work in psychology and attempts to extract algorithms that might capture some of the infant’s impressive cognitive growth.

One of the most comprehensive efforts at computer based modeling of early development following a Piagetian approach has been that of Drescher (Drescher, 1991). This used the concept of sensory-motor *schemas* drawn from Piaget’s conception of schemas in human activity (Piaget, 1973) and had similarities with early schema models by Becker (Becker, 1973). Unfortunately, Drescher’s implementation was a simulation and so many issues that concern embodiment were not exposed. Maes showed how Drescher’s approach can be improved by using focus of attention mechanisms, specifically using sensory selection and cognitive constraints (Foner & Maes, 1994).

A few models of infant grasping have been produced and some recent ones (Oztop, Bradley, & Arbib, 2004) hint that visual guidance may not be necessary for reaching, however none cover the growth of proprioception. There are many more models of sensory-motor coordination, and the vast majority of these 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, Guigon, & Burnod, 2001). Other models use basis functions (Pouget & 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. Interestingly, the model of Baraduc et al is one of the few that apply adaptation to the proprioception signals, and obtain good accuracy from very few input examples. While our models could be cast into a connectionist framework and, we believe, would produce identical results, we wish to formulate general methods for constraint models and so favour more explicit algorithms.

10. Conclusions

The implementation described in this paper builds sensory-motor schemas in terms of 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 system uses the LCAS approach to increase competence and two simple state variables are employed to detect saturation and trigger the lifting of constraints.

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.

The number of constraints actually used in the experiments was quite small, being less than the possible number discussed in Section 6 in order to limit the scale of the experiment. Nevertheless, the variation in behaviour was surprisingly rich and this suggests that we should see even more complex behaviour when all the constraints have been lifted; for reaching and grasping this will involve more motor components, a vision system, and even fully tactile sensing hands.

We notice that constraint lifting may be either a gradual or triggered process. For example, ignoring or taking account of a sensory channel is a binary decision but extending the range or probability of an action can be gradually increased. However, it is likely that even triggered stages overlap with previous stages and stages may blend rather than have sharp transitions. As an early researcher stated:

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

There is also the possibility that, due to unusual circumstances, behaviour might revert to an earlier pattern. This is known as regression and is also seen in our system when a previously unknown field is discovered during object touching behaviour. The stimulus of a new field is stronger than the contact event and the lapsed field-recording processes are activated.

Support for our approach comes from various data. For example, studies of the order of cell activation in the fetus report that the first cells to be detected as active are the somatosensory cells, then the auditory, then visual, and finally, the multisensory cells appear active (Meredith, Nemitz, & Stein, 1987). This suggests that there are advantages if proprioception leads in sensory development. Other work has also experimented with low resolution in sensors and motor systems and then shown that increasing resolution, leads to more effective learning (Gomez, 2004). Reduction in degrees of freedom obtained by staged development is also reported to be an effective strategy (Lungarella & Berthouze, 2002; Sporns & Edelman, 1993), as is the concept of constraints being beneficial to the emergence of stable patterns and helping to bootstrap later stages (Berthouze & Lungarella, 2004). Regarding our sensory-motor coordination method, we have avoided the long training times of connectionist methods and used a fast, incremental and constructive mechanism. This is in accord with several researchers who report that learning and adaptation can be very fast (Angulo-Kinzler et al., 2002; Rochat & Striano, 1999) and in some cases only one trial or experience is needed to alter behaviour. This is a new and interesting line of research.

We have argued that it is necessary to begin modelling development at the earliest possible behavioural stages. We agree that “early infant life is ... systematic exploration” (Rochat, 2003) and believe that robotics can learn much from infant psychology. Although most psychological theories are not fully articulated enough to allow testing via implementation, psychologists have built up considerable understanding and insights into cognitive development through experimental studies. We should make more use of this in autonomous systems research so that we might make steps towards the goal of “continuous development”, in the sense of (Prince et al., 2005). In the longer term, we hope this will lead to new methodologies for engineering robot systems and better models of human behaviour and growth.

Acknowledgments

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

- Angulo-Kinzler, R., Ulrich, B., & Thelen, E. (2002). Three-month-old infants can select specific leg motor solutions. *Motor Control*, 6(1), 52–68.
- Baraduc, P., Guigon, E., & 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 R. C. Schank & K. M. Colby (Eds.), *Computer models of thought and language* (p. 396–434). San Francisco, USA: W. H. Freeman and Company.
- Bernstein, N. (1967). *The coordination and regulation of movement*. Oxford: Pergamon Press.
- Berthouze, L., & Lungarella, M. (2004). Motor Skill Acquisition Under Environmental Perturbations: On the Necessity of Alternate Freezing and Freeing of Degrees of Freedom. *Adaptive Behavior*, 12(1), 47–64.
- Bruner, J. (1990). *Acts of meaning*. Cambridge, MA: Harvard University Press.
- Campbell, R., & Bickhard, M. (1992). Types of Constraints on Development: An Interactivist Approach. *Developmental Review*, 12(3), 311–338.
- Clifton, R., Muir, D., Ashmead, D., & Clarkson, M. (1993). Is visually guided reaching in early infancy a myth? *Child Development*, 64(4), 1099–1110.
- Drescher, G. (1991). *Made up minds: a constructivist approach to artificial intelligence*. Cambridge, MA: MIT Press.
- Foner, L., & 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*. NY: John Wiley.
- Gomez, G. (2004). Simulating development in a real robot. In *Proceedings of the 4th int. workshop on epigenetic robotics* (p. 1–24).
- Hainline, L. (1998). How the visual system develops. In A. Slater (Ed.), *Perceptual development: visual, auditory, and speech perception in infancy* (p. 5–50). Psychology Press.
- Hendriks-Jensen, H. (1996). *Catching ourselves in the act*. Cambridge, MA: MIT Press.
- Kalaska, J. (1995). Reaching movements: Implications of connectionist models. In M. A. Arbib (Ed.), *The handbook of brain theory and neural networks* (p. 788–793). MIT Press.
- Kalnins, I., & Bruner, J. (1973). The coordination of visual observation and instrumental behavior in early infancy. *Perception*, 2(3), 307–14.
- Keil, F. (1990). Constraints on constraints: Surveying the epigenetic landscape. *Cognitive Science*, 14(4), 135–168.
- Lee, M. H., & 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., & Chao, F. (2006). A content-neutral approach for sensory-motor learning in developmental robotics. In *Proceedings of the 6th int. workshop on epigenetic robotics* (p. 55–62). Paris.
- Lungarella, M., & Berthouze, L. (2002). Adaptivity through physical immaturity. In *Proceedings of the second international workshop on epigenetic robotics* (Vol. 94, p. 79–86). Lund University cognitive studies.
- Lungarella, M., Metta, G., Pfeifer, R., & Sandini, G. (2003). Developmental robotics: a survey. *Connection Science*, 15(4), 151–190.
- Meng, Q., & Lee, M. H. (2005). Novelty and habituation: the driving forces in early stage learning for developmental robotics. In M. E. S. Wermter G. Palm (Ed.), *Biomimetic neural learning for intelligent robotics* (p. 315–333). Springer-Verlag.
- Meng, Q., & 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)
- Meredith, M., Nemitz, J., & Stein, B. (1987). Determinants of multisensory integration in superior colliculus neurons. I. Temporal factors. *Journal of Neuroscience*, 7(10), 3215–3229.
- Oztop, E., Bradley, N., & Arbib, M. (2004). Infant grasp learning: a computational model. *Experimental Brain Research*, 158, 480–503.

- Pfeifer, R., & 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*. London: Paladin.
- Piek, J., & Carman, R. (1994). Developmental profiles of spontaneous movements in infants. *Early Human Development*, 39, 109-126.
- Pouget, A., & Snyder, L. (2000). Computational approaches to sensorimotor transformations. *Nature neuroscience*, 3, 1192-1198.
- Prince, C., Helder, N., & Hollich, G. (2005). Ongoing Emergence: A Core Concept in Epigenetic Robotics. In *Proceedings of the 5th int. workshop on epigenetic robotics* (p. 63-70). Lund University cognitive studies.
- Rochat, P. (2003). Five levels of self-awareness as they unfold early in life. *Consciousness and Cognition*, 12(4), 717-731.
- Rochat, P., & Striano, T. (1999). Emerging self-exploration by 2-month-old infants. *Developmental Science*, 2(2), 206-218.
- Rutkowska, J. (1994). Scaling up sensorimotor systems: Constraints from human infancy. *Adaptive Behaviour*, 2, 349-373.
- Sporns, O., & Edelman, G. (1993). Solving Bernstein's Problem: A Proposal for the Development of Coordinated Movement by Selection. *Child Development*, 64(4), 960-981.
- Thelen, E., & Whitmyer, V. (2005). Using Dynamic Field Theory to Conceptualize the Interface of Perception, Cognition, and Action. In J. Lockman & J. Rieser (Eds.), *Action as an organizer of learning and development, minnesota symposium on child psychology* (Vol. 33, p. 243-277).
- Tronick, E. (1972). Stimulus control and the growth of the infant's effective visual field. *Perception and Psychophysics*, 11(5), 373 - 376.
- Turing, A. M. (1950). Computing machinery and intelligence. *Mind*, 59(236), 433-460.
- Westermann, G., & Mareschal, D. (2004). From parts to wholes: Mechanisms of development in infant visual object processing. *Infancy*, 5(2), 131-151.

Figure 1:

Figure 2:

Sensory Space (S_1 S_2), Cartesian Encoding

Motor Drive Space (D_1 D_2)

Figure 3:

Figure 4:

Figure 5:

Figure 6:

Figure 7:

Figure 8:

Figure 9:

List of captions for figures

- 1: The experimental robot arm system
- 2: A plan view of the arm spatial configuration.
- 3: Sensory-motor coordination through 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.
- 4: Arm movements with no contact sensing. Actions are shown as vectors between start and end points in motor space.
- 5: Fields generated during the non-contact stage. Shown in Cartesian encoding of proprioception space.
- 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.
- 7: Fields generated after early object contact. Shown in Cartesian encoding.
- 8: Rates of growth in maps. Only initial field visits are counted.
- 9: Cyclic patterns of behaviour produced by several objects.
- 10: Transitions between three maps of different scale. Only initial field visits are counted. The Mapping Scale plot indicates the switching points.

Figure 10: