

A Developmental Framework for Cumulative Learning Robots

Mark Lee, James Law, and Martin Hülse

Abstract Developmental psychology is the study of human cognitive growth. However there exists a huge gap between the psychologist's theories and knowledge of behaviour and our ability to implement developmental processes in autonomous agents. In this chapter we describe 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 described and illustrated with results from experiments. We discuss how this approach is grounded in the topics of sensory-motor abstraction, intrinsic motivation (as novelty), and staged learning, and our belief that robotics can learn much from infant psychology.

1 Introduction

Research in Artificial Intelligence has produced some very impressive results over recent years and many of these advances have been implemented in robotics projects. But despite this success the realisation of truly autonomous robots has proved elusive and remains a difficult and significant challenge.

Any really intelligent robot should be autonomous in the sense that it can operate in the real world and can *cope* with new experiences and events by drawing on its previous experiences and building up cognitive competences and skills similar to those seen in humans and animals. This requires qualities such as adaptation, learning and versatility. But robots, like humans and unlike other computational systems, are embedded in the real world and have to experience noisy, chaotic, unstructured environments. Hence, completely novel experiences are unavoidable

M. Lee (✉) · J. Law · M. Hülse
Department of Computer Science, Aberystwyth University, UK
e-mail: mhl@aber.ac.uk; jxl@aber.ac.uk

and this demands more than simple adaptation or the learning of stable events; indeed, new learning processes must emerge as conditions change and new events, environments and relationships are experienced.

It seems that such cognitive flexibility, as readily seen in humans and other altricial animal species, is dependent on a prolonged period of parental care, during which occurs some remarkable processes of structured growth generally known as “development”. Given the long history of research into learning, in both animals and machines, it is surprising that only recently has any real attention been given to the concept of development as an important factor for artificial autonomous systems. This is remarkable because psychologists and other scientists have studied development in great detail and the idea is not new or original. Indeed, the great computer pioneer 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)

Most work in bio-inspired robotics has drawn on the burgeoning growth of research in neuroscience and brain science generally. However, most brain science is concerned with understanding structure and function, with less emphasis on how those structures arise and how they are shaped by experience. But fortunately, developmental aspects are now gaining attention and the field of developmental robotics has now become established (Lungarella et al. 2003). This approach assumes a developmental framework that allows the gradual consolidation of coordination and competence, and emphasises the role of environmental and internal factors in shaping adaptation and behaviour (Prince et al. 2005).

The challenge from this viewpoint is in finding effective algorithms for processes that support the development of learning and adaptation. There exists a very significant lacuna between our psychological theories of development and our ability to implement working developmental algorithms in autonomous agents. In this chapter we explore this area by first examining the role of staged growth, constraints and infant learning (Sects. 2 and 3). In Sects. 4 and 5 we describe our approach for building developmental algorithms, and in Sect. 6 illustrate them with results from experiments. In Sects. 7 and 8 the role of novelty for motivating an active learning architecture is examined. Finally, in Sects. 9–11, we discuss the broader context and some key research challenges that emerge from this approach.

2 Developmental Stages

The scientific study of human cognitive growth is known as developmental psychology. Experiments on children and adults produce data which inform the production of theories that could explain the growth of skill and competence over time. Unlike neuroscience, psychology has no direct access to brain processes and instead attempts to derive inferences about possible internal mechanisms and events through observations of behaviour. Hence, patterns of behaviour and behavioural dynamics become the currency of experimental investigations. This level of indirectness allows much scope for interpretation and variety in 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 appears before another: looking; reaching; standing; walking, etc. 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. 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). Very briefly, Piaget defined four periods in an individual's life: beginning with the sensory-motor period (up to 2 years) during which infants are not fully capable of symbolic representation; then the Preoperational period (2–6 years) is characterised by egocentric behaviour; then follows the Concrete Operational period (6–12 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. Other psychologists, 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; Kalnins and Bruner 1973).

All developmental stages have vague boundaries, and overlap and merge with other stages. They also show considerable temporal variation between individuals. Consequently, there is 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 apparently 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 any structures acquired to support development are shaped by experience (Spelke 1998). This debate, also known as the Constructivism versus Evolutionism argument, still continues after many decades and has various implications for psychological theory. Nevertheless, all viewpoints recognise the existence of stages as manifestations of development and their role in the growth of cognition appears to be very significant.

2.1 *Early Infancy*

We believe that research into developmental algorithms for robotics should be firmly grounded in the sensory-motor 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 sensory-motor 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 a vital process for autonomous robots; and (4) it seems likely that sensory-motor coordination is a significant general principle of cognition (Pfeifer and Scheier 1997).

Furthermore, although the sensory-motor period covers 2 years, we believe it is essential to focus on the very beginnings of this period, before speech, before locomotion, and before other competences have become established. We are inspired by the first three months after birth, when control of the eyes, head and limbs is just emerging and growing. 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 cumulatively 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 for autonomous robotics to learn from this scenario.

For further support for this view, that we must start developmental learning at the very earliest stage possible, see Smith and Gasser (2005) where developmental psychologists argue from “six lessons from babies” that initial prematurity is crucial for the growth of embodied intelligence, in both babies and other agents.

3 The Importance of Constraints

All processes of adaptation and learning must have some form 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 restriction on sensing, action or cognition effectively reduces the complexity of the inputs and/or possible actions, thus reducing the task space and providing a constraining framework which shapes learning (Bruner 1990; Rutkowska 1994). In robotics, such constraints might be seen as sensory bandwidth reduction or the restriction of some degrees of freedom in motor actuation. The reduced task space can then be explored through active learning and its structure captured in the

learned representation. When a high level of competence at the current 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 around 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. See Pfeiffer and Bongard (2006) for expansion of this key topic. We notice that the anatomy of an infant changes markedly from birth, when we observe narrow shoulders and hips, large head as well as relatively short arms and legs.

Sensory-Motor: All sensors have their own limitations; usually these are specified in terms of accuracy, resolution, response rates and bandwidth. Motor systems also have similar characteristics with additional features for dynamic performance. Most changes in such sensory-motor characteristics can be linked to maturational growth.

Cognitive/Computational: Constraints on cognition take many forms, not only relating to speed but also relating to information content and structure. Many constraints that affect artificial neural systems are now known, and similar effects are being found in the brain (Casey et al. 2005). 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 (Johnson 1990). Both neural and endocrinal support systems will have effects as they mature, for example neurogenesis is still under way 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. If the constraints are carefully structured, especially by another agent to assist learning or adaptation, this is known as scaffolding (Bodrova and Leong 2006). Examples of

scaffolding are seen in parental care, education, and many other social situations, where interactions or tools generate patterns in the environment in order to direct attention or action towards a goal.

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

4 The LCAS Approach

It is important to state that we strive for general rather than task-specific mechanisms of development. Consequently we emphasise explicit, abstract models and try to avoid prestructured internal representations or assumptions about internal belief states or internal causal knowledge. In particular we avoid the early adoption of neural network models and other connectionist methods as these can be difficult to analyse and interpret (Gasser and Smith 1998; Lee et al. 2006). Such methods also often entail extensive training schedules which is counter to most of the empirical evidence that indicates learning and adaptation can be very fast, and in some cases require only one trial of experience (Angulo-Kinzler et al. 2002; Rochat and Striano 1999). Accordingly, we appreciate the “content-neutral” methodology of Thelen and Whitmyer (2005) and we try to follow a similar approach.

We view “constraint lifting” as a key mechanism for progression towards increasing competence. Transitions between developmental stages are related to the lifting of constraints, although the nature of such transitions is not fully understood. It seems that the enabling conditions for transitions must be related to internal global states, not local events, because local activity cannot capture cumulative experience. For example, assume novelty is a motivating driver, then if a novel event occurs (a new stimulus or behaviour) this will raise some local attention or excitation levels for that particular event. However, successive similar stimuli will be less excitatory, both in time and space, and eventually similar stimuli may be experienced for all spatiotemporal possibilities. At this point, such events are no longer novel and do not raise excitation levels significantly. A global state indicator (being the spatiotemporal sum of all local excitation) will thus reach a stable plateau when no novel events have been seen for a long time. Thus, high competence at a level is equivalent to all incoming experience matching expectations, 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, e.g. when growth changes have effectively ceased or when experience has become saturated. They can then signal the need to enter a

new level of learning by lifting a constraint (such as 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 (LCAS), 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, mechanisms for sensory-motor learning or adaptation are incorporated at this point. A set of global measures need to be established to monitor internal activity and some kind of intrinsic motivation must be provided to initiate action. We use a simple novelty function as the motivational driver.

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. The “Act” stage generates spontaneous and varying patterns of action so that 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.

This general description of LCAS above is not meant to be prescriptive in detail. Constraint lifting should not be triggered by explicit mechanisms but should be the emergent effect of the computational processes involved. Similarly, we do not advocate the programming of the stages by following an explicit constraint lifting schedule; such a method would defeat the very purpose of investigating development. But it is necessary to study the role of constraints and their relation to staged behaviour, and so constraint schedules need to be investigated in order to further our understanding of their relationships and interaction, and also to support the design of experiments.

Our approach to developmental robotics assumes that the robot’s task is to learn as much as possible and we try to avoid pre-programmed competencies. This is an experimental stance and does not represent a position on the empiricist/nativist spectrum in developmental psychology (Spelke 1998). We note that nativists need to explain the origins of any innate structures that they propose and any experiments on learning may shed light on this, just as much as they may support an empiricist stance.

4.1 Representations

In early infancy changes in behaviour can be discerned where initially spontaneous, apparently random, limb movements gradually become coordinated and organised. During this process the proprioceptive and kineasthetic space of the limbs becomes calibrated and correlated with the motor space of the actuation system of the muscles. The spontaneous movements would be primitive examples of active

learning, observed as “motor babbling” which is seen at many levels of behaviour. The very earliest examples of such proprioceptive behaviour probably occur in the womb (Ververs et al. 1998) and could provide the first sensory-motor correlations. We believe early correlation between proprioceptive space and motor space should be the foundation for building internal models of local egocentric space and thus form a substrate for future cross-modal skilled behaviours.

Interestingly, the same issue is found in robotics research where it is necessary to coordinate the differing spatial frameworks of the various sensory and motor systems. For example, coordinating the spatial frame of an eye system with the spatial structure of a hand/arm system requires cross-modal relations to be established and understood; in this case, image-based information needs to be related to the spatial coordinates accessible by a multi-degree-of-freedom mechanism. Even within a single modality there are often various correlations that have to be established, usually between sensors that relate to a particular motor system. For example, when moving an eye or a camera to a new fixation target it is necessary to relate the desired displacement on the image to the associated motor action that will achieve the target.


These concerns reflect on the problem of representation; how should these coordinations be implemented in a computer model so as to best capture the correlation relationships, while allowing that these must be learned and not programmed? In order to follow the methodology mentioned above, we have adopted a simple and general model that permits wide variation without being committed to a particular learning mechanism. Our model also has many neuromorphic features as it attempts to be compatible with knowledge from neuroscience.

4.2 A Sensory-Motor Mapping Model

We now describe the general idea behind our topological mapping method for the representation of sensory-motor events and coordinations. In particular, we show how such structures can form the key substrate for a developmental style of learning and can be fast, accurate and flexible.

In several previously reported experiments (Lee et al. 2006, 2007b), we have used two-dimensional maps with explicit links between corresponding sensory or motor values for the representation of sensory-motor spaces. Although three dimensions might seem appropriate for representing spatial events, we take inspiration from neuroscience, which shows that most areas of the brain are organised in topographical two-dimensional layers (Braitenberg and Schüz 1991; Mallot et al. 1990). This remarkable structural consistency has inspired our explorations of the potential and efficacy of two-dimensional structures. Also there is evidence to believe that the human spatial system separates depth (distance from the body) from lateral displacement (up-down, left-right). For example, when depth is an object feature being detected (along with colour, intensity, orientation, etc.) then this can be processed in parallel with a second feature; an unusual effect that does not

Fig. 1 Field overlap on regular *triangular* mesh, 10×10 fields, $r = 0.6$



apply to other features (Nakayama and Silverman 1986). This suggests that a 2.5D architecture is appropriate, with lateral and vertical locations as the axes of a 2D map, and depth values recorded at specific locations. A 2.5D structure is different from a full 3D form in that only one depth value can exist for each location; but this makes sense in an egocentric space where the nearest object occludes all those behind it.

A typical mapping will consist of a 2D array representing two sensory or motor variables, known as a *map* or *surface*, connected to another 2D array by a set of *links* that join points or small regions, known as *fields*, in each array. The links are bidirectional enabling access in both directions and are collectively known as a *mapping*. Fields are local regions of equivalence and are defined by a boundary function; we generally use simple circular fields. With circular fields it is not possible to completely cover a surface without some overlap between fields, unlike with pixels. For a uniform distribution of fields on a 2D grid, the most efficient surface covering is given for an equilateral triangular grid, where the minimum radii to ensure complete covering is 0.577 (with unity grid spacing), see Fig. 1. For this case the area of overlap is only 21 %, (i.e. for any field 79 % of its area is not shared with another field).

4.2.1 Field Distributions and Overlap

Overlapping structures may seem inefficient or inappropriate for computational implementation but it is well known that overlap occurs in many neuronal and sensing mechanisms in the brain (Carreira-Perpinan et al. 2005). For example, in the eye the sensing receptors do not physically overlap, but they are connected to ganglion cells that provide functional overlap (Sterling 1999). We believe this property is responsible for some interesting and very useful effects. We note that very large overlaps can have little purpose, because, in any given locality, many of the fields will have almost identical coverage. Thus a stimulus would give the same

effect over a wide area. Two possible functions for highly overlapped fields might be redundancy and low pass filtering. However, lesser overlap, when the field radii are below the grid spacing distance, appears much more promising.

Our mapping method consists of two processes: the creation of new fields on a map surface; and the generation of explicit links between two fields on different maps. A field captures a local region on a map such that all stimuli within the field can be represented and transmitted to other surfaces through a single field reference point. For circular fields the reference is the centre point. Links between maps are established through correspondence; if a field in one map can be reliably observed to become excited in temporal correlation with another field on another map, then an explicit link is created between the fields' centres. This is equivalent to finding the strongest weights between elements in a fully connected 2-layer network after training for temporal correspondence. Using this criteria gives very strong links for just a few fields and no coordination at all for all the others.

This simple Hebbian learning system (Martinetz 1993) has proved very effective in building maps in systems with strong and stable correlations. A generalisation of this approach is to use probability density functions to represent the unknown map linkages during learning. Variations of this include mechanisms based on radial basis functions (Pouget and Snyder 2000). This mapping method, although simple in concept, has several valuable properties. The number of links needed to effectively map between two surfaces can be quite low and increases with desired accuracy. This means that learning such maps can be very fast; we have grown mappings on laboratory hand/eye robot systems and achieved complete mapping coverage of a reach space of 3944 cm² in real time in 5 h (Hülse et al. 2010b). This scheme can also be adjusted so that links may be removed or changed when errors are detected thus providing adaptation and placticity.


5 Building Maps from Experience

Questions now arise as to when fields should be created and how they should be organised on a mapping surface. This concerns both the initial placement of the fields on the map and methods for exploring the sensory-motor spaces so that fields are populated with data from experience. In addition, we also need to consider the possibilities for adaptation of developed maps to take account of external changes.

5.1 Field Generation

There are essentially two main parameters for configuring the fields that comprise a map: they could either be of varying or fixed size, and they could be generated

Fig. 2 Fields being generated from a polar grid



onto a grid (prior structure) or they could be placed at any stimulus location (free format). This gives four possibilities in all, each of which we now consider in turn.

Assume that a learning process generates stimulus points, $p_i = (x_i, y_i)$, on a two-dimensional surface, S , which is initially empty. If a stimulus point is already covered by a field on S , i.e. is within the radius of some existing field, then no action is required. But if p_i is not covered, then a new field must be generated for this location.

Uniform size fields on a regular grid For this case, over the map S we arrange a grid of fields with uniform spacing of centre points. Initially all fields are unassigned but an uncovered stimulus point will cause the nearest field to be assigned to the map. Eventually all points will be covered by one or more fields. If the grid of fields has low levels of overlap ($r <$ the grid spacing), for example as in Fig. 1, then there will always be places covered by only one field and so eventually all fields in the grid will be used. Conversely, with large overlap many of the possible fields will not need to be generated, as each point on the surface only needs to be covered once. Figure 2 shows a retina design with considerable overlap but, during learning, fields are only taken from the grid as needed and the covering process stops when every possible stimulus point has been covered. It may seem unprincipled to use a pre-structured grid but topographic maps are widespread in the brain. There is evidence that this topographic structure is determined during neurogenesis by many influences and *both* genetic and experiential inputs have strong effects (Goodhill and Xu 2005). It is possible that genetically encoded neural growth patterns provide

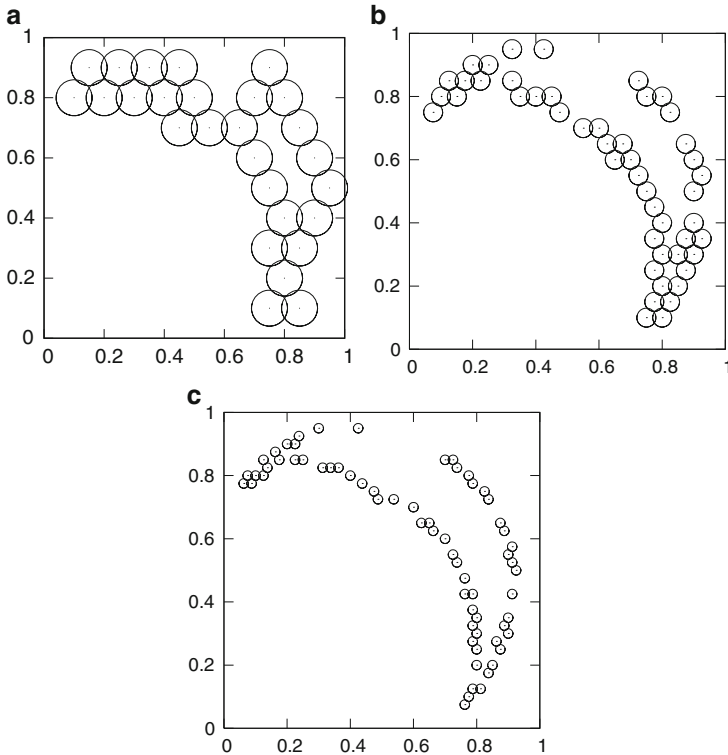


Fig. 3 A single map composed of three layers, each of different field sizes. (a) Large fields. (b) Medium fields. (c) Small fields

regular arrays of neural sheets and then the interconnections are established by a separate process of coordination. There is also evidence that neurons can expand their receptive fields in order to adapt to a damaged area caused by a lesion (Einarsdottir et al. 2007) and we note that such plasticity is better served by a uniform grid structure rather than an irregular covering of fields.

Variable sized fields on a regular grid This case is the same as above except that the field sizes are individually set according to the space available in their local neighbourhood (Meng and Lee 2007). For experimental purposes we found it more insightful to maintain several uniform maps, each with a different size of field. Figure 3 shows three such surfaces from our experiments. When a new stimulus point requires a new field we processed all the maps in parallel and so the three maps are built up together. When later using the maps to select a field, the large fields can be used for crude but fast coverage of the space, while the smaller fields require more learning but give finer articulation. This provides options for accuracy/speed trade-offs as and when needed. This effect was also reported by Gomez et al. (2004).




Fig. 4 Aligned neural fields of different sizes in the superior colliculus, from [Wurtz and Goldberg \(1972\)](#)

It is interesting that similar layers of fields of different sizes all responding to the same stimuli location have been found in the superior colliculus and elsewhere in the brain ([Wurtz and Goldberg 1972](#)). Figure 4 is taken from [Wurtz and Goldberg \(1972\)](#) and reveals increasingly larger field sizes at deeper layers of the superior colliculus. Note how a stimulus point in field 1 is covered by all fields, with each successive field usually overlapping all of the preceding field.

Uniform size fields with irregular locations In this case fields are simply created with their centres located at the incoming stimulus points. This is generation-upon-demand and so the shape of the final coverage will not only be irregular but will also vary considerably depending upon the order in which the stimuli arrive. It may not be appropriate to be driven by event order, as very idiosyncratic mappings can develop which might not easily generalise over different tasks. However, this objection can be overcome by allowing fields to be replaced by some form of decay mechanism that removes infrequently used fields.

We have found this to be a simple and effective generation method for situations where there are no topographical requirements or when the nature of the spaces is unknown. For example, in correlating a robot arm with a motorised camera system we found that a mapping of arm end-point with camera gaze point gave a very effective coordination scheme ([Hülse et al. 2010b](#)). Figure 5 shows such a mapping. This system was driven by an exploratory algorithm that looked for unmapped areas and attempted to find new points equidistant from the nearest existing fields.




Fig. 5 Mapping between robot arm (*left*) and eye system (*right*). Fixed field size, free location. Similar colours indicate linked fields

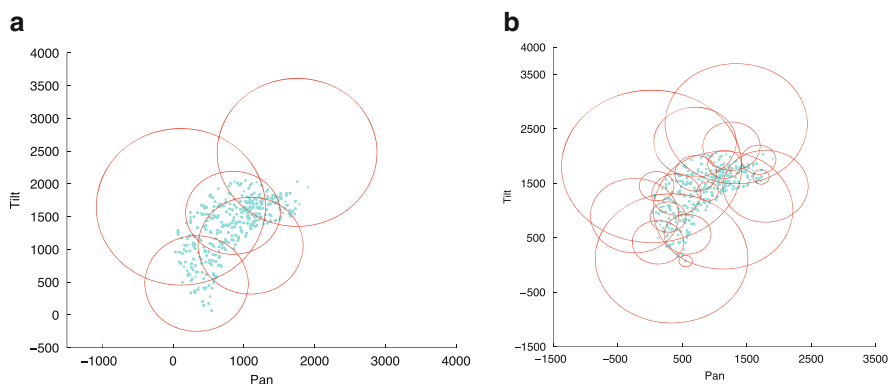


Fig. 6 Variable field sizes and locations. (a) Coarse coverage. (b) Finer coverage

Variable sized fields with irregular locations Completely irregular field growth is also possible. We have experimented with mechanisms where any uncovered stimulus causes a new field to be created that is precisely centred on the stimulus, with the size of the field determined to maximally fill the gap between the nearby fields. Figure 6 shows some results from radial basis function experiments (Meng and Lee 2007). The left diagram shows how only five fields initially developed to cover a set of points; crudely but quickly. The right-hand diagram is after further learning where more, and smaller, fields give a finer coverage with more resolution. A similar reduction in cortical field areas relating to infant growth in object recognition ability has been reported (Westermann and Mareschal 2004).

We notice that grid placement is useful for situations where some external structure must be taken into account. This is seen in sensory maps that must reflect the structure of the sensing system. For example, peripheral vision has a much lower resolution than that of the foveal region and so we use increasingly larger sized fields towards the periphery in retinal maps. By designing a grid before use, we can impose desired constraints on field size and placement.

5.2 *Populating Fields Through Exploration*

The next issue is how map building processes are to be driven and organised in a learning or growth scenario. Consider the original problem at the point where no known structure or patterns have been discovered. Hence, no links yet exist. Assume we have two 2D surfaces, S and M , and we wish to establish how they are related. The surfaces both have two variables: (x_i, y_i) for S and (k_j, l_j) for M ; and we can perturb these in order to detect any correlations. However, it is generally not possible to vary sensory inputs at will (at least not directly) and so only M is available for exploration. This means we can vary M and observe any effect on S ; we cannot operate in the other direction. Thus, starting from a condition of no prior knowledge, some initial motor action must be performed, and the least specific action is simply to start moving in some arbitrary direction. Such action will eventually terminate when the physical or anatomical limits are reached for that particular system. If this is followed by similar movements in other directions, then the effect is produced of exercising M over its range of variables. This behaviour will explore the maximum and minimum extent of M over its range limits. We note that this appears very similar to the behaviour known as motor babbling in human infants (Piek 2002; Piek and Carman 1994). If the two variables for M are independent, then a rectangular plot will emerge for the boundary, showing the ranges of k_j and l_j along the axes. Now, the values of x_i and y_i may vary in a complex way with M but if the motor values are constrained to their extrema, then the S values will similarly describe the limits of their range. The plots thus produced will display the boundaries of the mapping for the operating regions of S and M . This assumption rests on the surfaces being smooth and continuous and having planar topologies. Thus, no points can exist *external* to the S boundary: otherwise, some regions of S would map into 2 or more separate places in M ; which could not provide a mapping. Figure 7 shows the (kinesthetic) sensory space produced for a two-limb jointed arm as the motors drive the angles of the joints through their extremities. Arrows are shown on the figure to illustrate the effect of sweeping one motor variable at a time. It is important to notice that such a structured strategy is not necessary for discovery of the boundaries. It is significant that *any* motor action that stops at an extrema for one of the variables will find a boundary point and the corresponding point on S will be revealed.

It is noticeable that while the motor surface covers the full range of the variables, as would be expected if the system is exercised over its extent, the sensory fields do not cover all possible locations on their surface. This is a common effect as sensory systems will often cover larger spaces than those of associated motor systems. For example, a camera on a pan-and-tilt head may have a wider field of view than the pan limits; and the human eye can see more than the eyeball can turn to. This means that the off-boundary (*internal*) fields in the motor surface will be mapped across to the *internal* region of the sensory map and the sensory fields must be effectively distorted to fit *within* the boundary shape. Figure 8 shows three different sensory examples for the motor pattern in the lower left. The motor system of the arm has two independent motors each with limits on their extent. This gives a rectangular




Fig. 7 Boundary fields can be established first

shape for the fields experienced at those limits (lower left in Fig. 8). The other three plots show the corresponding fields for three different sensory configurations (using different geometric configurations of the proprioceptive sensors).

A consequence of this kind of motor babbling behaviour is that the extremities of the space will be explored before the fine detail of the internal regions. If we consider using a regular grid for the fields, as in Sect. 5.1, then the boundary distortion can be seen as a kind of warping of M to S . Consequently, we could warp the regular grid to match the boundary and thus obtain a good first estimate of the locations of the internal fields *even before any such data points have been experienced*. Standard image warping methods are not useful here as they use parametric transformations that do not handle local distortions well. Various kinds of local distortion methods do exist but a far simpler computational approach is provided by the elastic membrane concept. An elastic sheet can be pulled and stretched in various directions to fit both local distortions and global warpings. The method uses the idea that elasticity is a simple relation between distance and force whereby each molecule of material moves to a position that minimises the total of the forces from its neighbours. So an effective relaxation algorithm is easily implemented by minimising the sum of the distances between each node in the grid and its immediate neighbours. This elastic neighbourhood consists of the six nearest fields in a triangular array (or four or




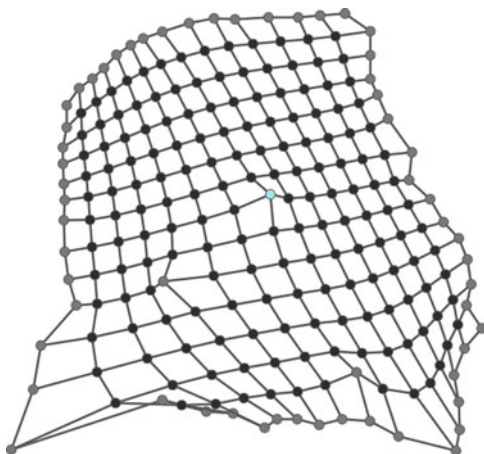
Fig. 8 Different boundaries for different arrangements of (muscle-based) sensor systems. *Bottom left*: motor map, *Top left*: Cartesian space (ideal) as would be seen by overhead camera, *Bottom right*: space of arm vector from hand to shoulder, *Top right*: space of vector from hand to body-centre

eight for a rectangular system) and as a node is moved to a new position so the local neighbours are pushed and pulled in the same direction but to lesser amounts the further away they are. Figure 9 gives an example of a grid with both boundary distortion and some internal node displacement. The algorithm calculates the error in each node and relaxes over all nodes until the total error in the system falls below a threshold. If we use the elastic sheet method to fit all the (originally rectangular) edge nodes in S to the new boundary (matching corresponding fields), then we will obtain a warp of the regular grid that suggests reasonable estimates for the locations of the internal fields.

5.3 Adaptation and Plasticity

Although we expect mappings at the lower levels of sensory-motor experience to be relatively stable, the possibility of error is always present and any mapping might be

Fig. 9 Boundary and internal distortions of elastic grid



required to adjust its structure in some way. This could be a local or a global effect: either a relatively small number of fields may be found to be located in error and need adjustment or the whole map might be in error and then a full scale remapping is needed.

Local Adaptation

Errors in field placement may occur due to noise and tolerance effects in sensing and/or motor structures. For example, an action may not be exactly repeatable due to low muscle control or environmental disturbance. These conditions can give rise to local error and the need for local adjustment.

The elastic sheet method, described in the last section, allows any point to be adjusted at any time, with the neighbours also making compensatory adjustments. This is ideally suited for the correction of local errors, where a field centre is found to be incorrect and is to be moved. Thus, if a local variation is created by some sensory or bodily change, then the discovered effects can be used to adapt the mapping by either small movements of the fields concerned or reassignment of a few links between the relevant maps.

Re-calibration and Realignment

In cases of major reconfiguration of sensory or motor systems the adaptation required may be so severe that complete realignment of all of a mapping may be required. This can happen when, for example, a camera in a hand/eye system is shifted to a new position relative to the rest of the system. In humans, similar disruptions are experienced in prism adaptation experiments where prisms applied to the eyes cause major and total redistribution of the visual image ([Redding and](#)

Wallace 2006). Such global changes will essentially involve all the links in a mapping being reassigned to quite different (i.e. non-local) fields. This is a serious disruption to any agent and presents a major re-learning challenge. Fortunately, such situations are generally rare and dealing with these cases efficiently may not be so important as ensuring good local adaptation.

The ability to change or remove links in error provides a degree of plasticity in the mapping that covers these problems. We have experimented with a mechanism for adaptation that allows links to be removed if they become unused (Hülse et al. 2010b). We gave every link an “age” value. When a link is first established and whenever it is subsequently used successfully then its age is set to zero. Otherwise the age of every link is increased at each application of the mapping. Any links that are not found useful are deleted from the mapping and this allows space for new links to be established. Thus, small local adaptations may occur continuously. See Hülse et al. (2010b) for further details. We have found that a threshold for the minimum age of links effectively controls the total number of links in a mapping and is a very useful parameter in experimental explorations of map growth and development.

6 Examples of Emergent Behavioural Stages from LCAS Experiments

In this section we present a few examples of our method to illustrate the concepts and mechanisms involved. As an experimental design we explored various sensory-motor subsystems individually. This approach also fits well with the idea of early constraints being quite severe and apparently limiting activity to one or two modalities at a time. Consequently, we started with an empty system and considered what sensory or motor components should or could be developed first. The eye is particularly active after birth, followed by limb movements, and then hand/eye interaction. Accordingly, we carried out separate implementations on eye saccading (Chao et al. 2010) and arm reaching (Lee and Meng 2005), and then combined these in hand/eye experiments (Chao 2009).

Consideration of the motor system shows that the kinaesthetic sense plays an important role in motor control and in spatial cognition. For example, when a visual stimulus is to be brought to the foveal region of the retina it is not just the sensed image and the eyeball muscles that are involved: the proprioceptive sensors in the muscle spindles provide very important data on eye position, or gaze. We found that proprioception was very significant in all the systems we investigated.

Our first study on the eye implemented a simple mechanism for saccade learning. We assumed that no calibration between image and eyeball position could be achieved before images were available (i.e. before birth) and used the methods described above to generate active exploratory behaviour and build mappings between stimuli fields on the retina (image) and motor positions of the eyeball

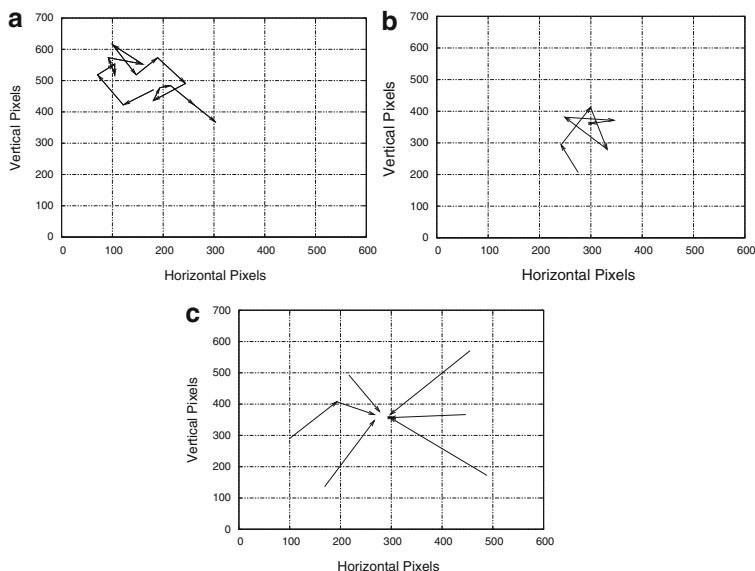


Fig. 10 Qualitatively different behaviours during learning. (a) Early stage trace. (b) Intermediate stage. (c) Final stage trace

(camera gaze direction). The aim was to learn a mapping between image locations and motor values so that a single direct movement (saccade) could be made to bring any image point to the centre of the image (fovea).

At first, when no fields or links had been established, the eye movements appeared as random walks, eventually finding the centre region. After a few fields had been created, the random moves would often find one of these and then move among the fields nearer the centre. Finally, single saccades would appear as the map became fully populated. We were surprised to see that the results showed clear qualitative differences in the behavior as the maps grew.

Figure 10 shows example traces clearly exhibiting different stages. The first trace performs 15 movements before finding the target, the second trace required 6 moves, and the last trace shows 6 different saccades (superimposed), all but one being single motor acts. The results were classified into three types: no local field exists near the stimulus; a neighbouring field is found and used; and a stimulus-covering field exists. These were plotted in Fig. 11, which shows three runs for each type to illustrate the variation. The labels correspond to those in Fig. 10. Type A is seen to start first but reaches a plateau after about 60 moves, while type B starts later and also plateaus later. The single move or fully learned saccades do not appear at all during the first 18 moves but grow fast until they become the only type to exist. Because fields are being generated for every movement the map builds quickly; in fact, Fig. 2, on page 187, shows the retinal map for this experiment after 94 fields have been entered. Already, much of the total space has been covered, after only around 180 learning attempts.

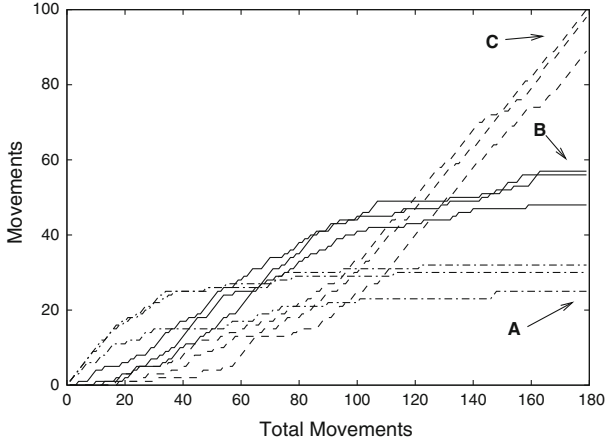


Fig. 11 The growth rates of three types of saccade behaviour




Fig. 12 Exploratory arm motions shown as vectors in arm joint space

Our next experiment imposed the constraint that the vision system would be inactive while the arm was active. This is equivalent to limb movement out of sight of the eye. Again we started with an empty map system. A rest position for the (single) arm was provided, equivalent to a low-energy state, and spontaneous motor values were set to drive the arm through the workspace. Eventually the arm comes to a halt, an unexpected event, and fields are generated. The process is repeated and the fields relating to the boundary of its workspace are soon discovered, as shown and described in Sect. 5.2. Figure 12 shows traces of motor acts; the darker section in the middle are movements from the rest position to the fully extended arm position (these were the initial motor settings) and other traces can be seen as the arm makes spontaneous moves and reaches various locations on its operating boundary.

As the map becomes populated so the rate of field discovery saturates, as does the opportunity to learn. To enable further learning, a constraint can be lifted, and in this case, we activated a tactile sensor in the robot end-effector. This allowed the arm to make contact with objects, interrupt the action, and record a new spatial sensation. As the arm touched objects in different locations so the internal fields in the map were created and eventually the map was completed. After this point any target “felt” location could be reached by a single direct arm movement.

The final constraint to be lifted was to allow both arm and eye to operate together. As both had near fully complete maps from the previous stages, this stage involved the creation of new mappings that relate the visual space of the eye to the reach space of the arm. We found the visual gaze space (i.e. the angles of eye fixation) to be the most appropriate frame for the integration of the two subsystems (Chao 2009).

It is interesting to notice the kinds of behaviour produced from this series of experiments. We observed a progression of qualitatively distinct behavioural patterns:

Visual stimuli fixation - through three stages from eye wandering to direct saccading, during the creation of the image map (constraints: eye only active)

“Blind groping” - actions mainly directed towards the body area (constraints: arm only active)

Extended groping - directed towards the bounding limits of the agent’s egocentric space, during the creation of the boundary space map through arm proprioception (constraints: arm only active)

Unaware contact - seen as pushing or ejecting objects out of the local environment, due to a constraint on the tactile sense (constraints: arm only active)

Contact sensitivity - or “sensitive groping” where limb movements are interrupted by tactile sensing events, and the non-boundary space map is constructed (constraints: arm and tactile active)

Repeated cycles of contact - observed as repeated “touching” behaviours directed at detected objects (constraints: arm and tactile active)

Hand regard - or hand fixation, when the eye sees the hand as an object but one whose movement correlates with arm activity. The “object” is marked as a special case (constraints: eye and arm active)

Arm follows eye fixations - when the location of a stimulus, found by the eye, is mapped into the arm space and excites arm action to reach to the same place. This is the basis of reaching and grasping of seen objects (constraints: eye and arm active)

Eye follows tactile contact - the converse case, where the arm has touched a stimulating object and its location is mapped to the eye system which then saccades to fixate on the object (constraints: eye, arm and tactile active)

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

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 et al. (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 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 appears to be supported by our 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 the structure of tactile sensing data is explored. 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. From these considerations it is clear that both components of the haptic system could develop together or proprioception could lead tactile and somatic sensing; but tactile cannot lead proprioception. On reflection we see that this is a logical necessity because the tactile system must rely on an existing spatial frame if its experiences are to have any spatial context or meaning.

Regarding environmental constraints, we have only used 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 from the agent’s personal space. 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.

The size of the fields is a useful constraint that is easily overlooked. If large fields are generated initially, then a rapid but crude mapping of space can be obtained. When this is no longer creating new experiences, then the field size can be reduced thus refining the accuracy with new map entries. 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](#)).

7 Novelty as Motivation

Motivation is an essential function to drive autonomous development, with current approaches often using externally driven processes (e.g. [Bullock and Grossberg 1988](#); [Caligiore et al. 2008](#); [Gasser and Smith 1998](#)). However, internal drives are necessary for true autonomy and we employ novelty-based functions as intrinsic motivators to drive autonomous development through increasing complexity. Examples of similar approaches and effective results are seen in [Kaplan and Oudeyer \(2003, 2007\)](#) and [Schmidhuber \(1990\)](#).

In our work we have used a simple novelty indicator to motivate learning and trigger the removal of constraints. When the robot encounters a novel stimulus it will repeat the action that caused that stimulus to occur. As the action is repeated, the novelty dissipates, and the robot has an increasing tendency to perform alternative actions. Over time the robot finds and performs cumulatively more of the actions available to it, and so the number of remaining novel actions diminishes ([Lee and Meng 2005](#)). As fewer novel actions are found, the rate of learning in the robot saturates. In order to enable learning to continue, a constraint can be lifted, as described in Sect. 6. This opens up a new range of actions for the robot and provides a new source of novel stimuli to investigate.

To implement this idea we allow the fields to store several variables; this is equivalent to maintaining an interlayer for each variable in a map surface. These variables can include stimulus-type quantities for the field location (e.g. depth, colour, intensity), and also excitation, activity, and any other indicators that prove useful as a substrate to support learning (e.g. object markers for short term memory). We use an excitation value to record the current salience or importance of a field's contents and an activity value for the usage the field has received. Novel events at a given spatial location bestow high excitation on the relevant field. Excitation levels gradually decay with time and are also reduced by habituation following repeated stimulations. Thus a few constants are required for habituation rate, recovery time, excitatory decay and possibly other influences. This mechanism allows a simple selection function; the field with the highest excitation provides the motivation for the next action.

Notice that we do not record novelty as a value—rather novelty increases excitation and in this way the meaning of novelty can change. Thus the effect of different events on excitation will change with experience of those events. Hence this approach covers a range of new events: new fields, new coordinations (links), new action changes, new sensing stimuli, new cross-modal events, etc. Global excitation (the normalised sum of all the individual excitation levels above a threshold) is an indicator that can signal low attention and can trigger a return to spontaneous motor action. Global activity (the normalised and inverted summation

of the field activity levels) decreases with familiarity and can indicate saturation. We have previously suggested that the degree of motor noise in an action is related to muscle tone and that tone should increase with excitation. If this is implemented, then highly focussed action (on targets of high interest or novelty) will have less noise and better accuracy, while low excitation will accompany low tone and higher motor noise resulting in more exploratory action.

8 Developmental Action Formation

In Sect. 2 we argued for the importance of grounding research in developmental robotics in the sensory-motor period of human infant development. However, developmental psychology supports the theory that early motor skills are related to perceptive and cognitive development (Thelen 1995), which varies from child to child. Although there is no single development sequence for all infants—they develop in their own ways and at their own speeds with some variation—they do generally conform to a common set of developmental stages. Through our research, we have identified some of these common sequences, which provide the foundations for our work on the developmental formation of actions.

From the literature on child development we have constructed a large, general time-line chronicling the observable development in infant sensor and motor behaviour over the first 12 months. The information is too extensive to cover here, but brief summaries can be found in Law et al. (2011) and Hülse et al. (2010a).

Based on this time-line we have generated some general sequences of development appropriate for implementation on humanoid robots. For example, Fig. 13 shows a partial development sequence for the upper body of a humanoid robot. Note that this sequence only focuses on motor development; similar charts cover sensory growth and attentional preference. Development of motor function is indicated by shaded areas, which get progressively darker as control improves (note that the termination of a shaded bar does not indicate the termination of that skill, but that it has been sufficiently developed as to no longer appear in the literature). It is interesting to note that the sequence for motor development begins with the eyes and progressively moves down the body, through the neck and arms, with torso and wrist control being refined last. The exception to this is the hands, which become active in the first month, but continue to develop until the tenth month. This time-line, along with its counterparts, begins to address the problem of how constraints should be ordered in a robot. For other work on time-lines, see Metta et al. (2009).

The framework to support this emergence of skill is provided for in our systems by an arrangement of constraints. Table 1 shows how, following our developmental time-lines, a series of constraints can be assembled to support the emergence of actions. In this case, we suggest a possible sequence for developing gaze-directed reaching beginning with the learning of arm proprioception and eye saccades.

Activated abilities are represented by crosses in Table 1. This does not necessarily mean that the corresponding function is completely unconstrained, but that it is

Table 1 A possible framework for staged development of integrated reaching

Development stage	Eye		VM			Arm			Grips			Observed behaviour
	M	P	LR	HR	IOR	Ex	M	P	Pa	Pi	Cu	
1 Eye saccade	x	x	x									Eye can saccade to fixate on any stimulus on the retina
2 Inhibited visual search	x	x	x		x							Ignores recurrent visual stimuli
3 Hand-eye coordination	x	x	x			x	x					Reaches to fixation point
4 Reach guided eye-saccade	x	x	x		x	x	x					Fixates on reach point
5 Self-aware reach	x	x	x		x	x	x					Ignores own hand
6 Retaining grasp						x	x	x		x		Holds objects with reflex grasp
7 Intentional gross pick-up	x	x	x			x	x	x		x		Reaches towards objects and picks them up with palmar grasp
8 Intentional precision pick-up	x	x	x			x	x		x	x		Reaches towards objects and picks them up with pincer grasp
9 Selective grasping (combination of tactile and visual data)	x	x	x		x	x	x	x	x	x		Reaches towards objects and picks them up with the appropriate grasp
10 Visual search with identification and memorization	x	x	x		x							Attraction to specific objects and refining of visual selection criteria
11 Tactile search with identification and memorization						x	x	x	x	x		Attraction to specific objects and refining of tactile selection criteria
12 Integrated object pick-up	x	x	x		x	x	x	x	x	x	x	

At each stage the system is constrained so that it only has access to data marked "x", where *M*, motor values; *P*, proprioceptive sensor values; *LR*, low resolution visual layer; *HR*, high resolution visual layer; *VM*, visual memory; *IOR*, inhibition of return; *Ex*, visual excitation; *Pa*, palmar grasp; *Pi*, pincer grasp; *Cu*, cutaneous sense

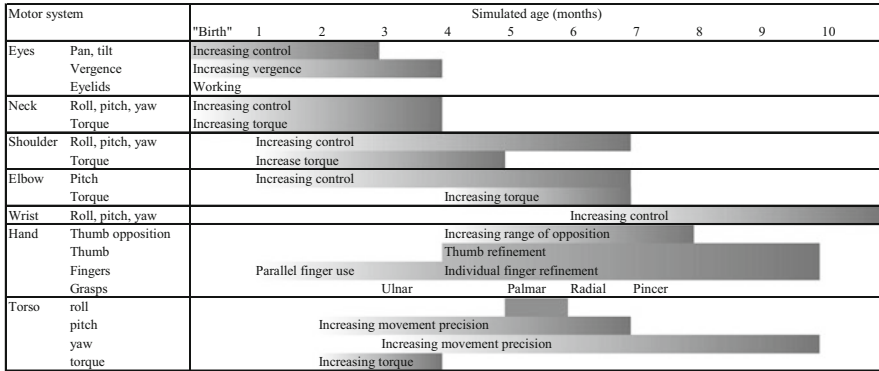


Fig. 13 Partial motor development sequence for a humanoid robot. *Shaded* regions relate to periods of development of each ability as observed in infants. *Darker* shading indicates more advanced ability

available in some form for the system to use. For example, when sensor and motor maps become available for use they may initially have additional constraints on their resolutions. In this example, the robot would start out at the first stage with access only to the eyeball proprioceptive sensors and motors, and low resolution visual feedback. This constrains the robot and focuses its attention on learning to saccade to simple stimuli. When sufficient learning has taken place development moves onto the second stage, where additional functionality is enabled. A new round of learning begins, this time reincorporating skills learnt in the previous stage. At each stage, one or more constraints on functionality are removed, allowing the robot to learn new skills. In this way, the robot progresses from being able only to make uncoordinated motion with an eye or an arm to being able to reach to a seen object in an integrated and goal directed behaviour.

It is important to note that the stages in Table 1 are not fixed in their order, nor are they necessarily triggered in isolation. Some stages may be learnt in parallel, and similar levels of development may appear in different orders. The robot is able to influence the order of constraint removal, so we would expect initially identical robots to develop in individual ways, much as human infants would. Importantly, the constraints do not directly control the developmental stages, but simply release more complexity to the learning processes. Thus stage transitions are emergent; their ordering and timing are not easily predictable. Indeed, the system may regress to earlier stages when an action cannot be successfully learned due to gaps in the system’s previous experience.

This raises the question of how different sensory-motor systems can be combined. Consider the example of three important maps: the image map from the retina of the eye; the gaze map which maps the orientation of the eye fixation point; and the reach map of an arm that records the places the hand or end-effector can occupy. These are all *spaces* in the sense that each one models the structure of a particular sensory-motor system and can relate an action to its effects. However they are quite




Fig. 14 Combining image, gaze and arm maps

different and distinct, yet must cooperate during behaviour. For example, if the eye notices a stimulus in a peripheral area of the retina, it must rotate the eyeball to bring the gaze to focus on the object; this involves a mapping from image map to gaze space. Next the arm can reach for the object but to do so the gaze location must be mapped into the reach space of the arm.

We have investigated several methods for combining gaze and reach mappings, such as those briefly mentioned in Sect. 6. In Hülse et al. (2009b) we used a robot arm to place objects within a camera’s field of view, and the system correlated the position of the object in the gaze space with the known positions of the arm joints. By repeating the process with objects placed at different locations, the system built up a mapping of gaze space to reach space. This system is outlined in Fig. 14.

Visual stimuli on the retina are mapped to eye motor movements in the mapping labelled “mapping for eye-saccade”. Here, a motor movement is mapped to a stimulated field in the retina map such that it will perform a saccade (as described in Sect. 6). This saccade results in a specific absolute motor configuration p , and the range of absolute motor configurations define the gaze space.

Points in the gaze space can be associated with physical locations in space by mapping them to arm configurations. In our experiments, the arm placed a coloured object within the field of view. The joint positions of the arm were then mapped to

the visual field stimulated by the object. This “mapping between reach space and gaze space” enables the robot to reach to a point of interest in the scene (Hülse et al. 2010b, 2009b).

Together these two mappings enable the robot to saccade to, and reach to, a target. However, to prevent the robot repeatedly saccading and reaching to the same targets, a visual memory is included, labeled VMGS (Visual Memory in Gaze Space). The VMGS stores the absolute motor configurations, p , of the active vision system resulting from a saccade. For each visual stimuli detected, the system is able to predict the corresponding gaze space configuration by applying the mapping for eye-saccade. By cross referencing these targets with those stored in the VMGS, the system can inhibit visual stimuli that have already been saccaded to (LVMM). The overlaid saliency map (OSM) contains the difference between the stimuli on the retina and the previously examined ones, thus enabling the system to only perform saccades to new stimuli (Hülse et al. 2009a).

9 Research Challenges

The work described above presents us with a number of challenges for the future. Some of these, and hints at solutions, are summarised below:

What field sizes should be used to create the various mappings? We have learned much about different field structures and how accuracy of representation is dependent on field size. But there is much more to be understood about the granularity and structure of human egocentric space. We suspect that further work on this topic might reveal some general principles for developmental robotics.

What are the properties of overlapping arrays? The idea of overlap is rather contrary to intuition and existing mathematical approaches. However, overlapping neural and sensory structures are ubiquitous in the biological world. Although overlap suggests engineering problems like signal crosstalk we believe it has considerable beneficial properties, especially for spatially grounded systems. Further theoretical work on models of topological maps is required, e.g. (Carreira-Perpinan et al. 2005).

How can mappings between multiple sensory-motor systems be combined? The work described here has focussed on creating mappings between one or two sensory or motor systems. Although these provide foundation building blocks, more complex activity requires the combination of multiple sensory-motor systems. Full scale humanoid robots will need many mappings across many different subsystems. This raises important organisational questions involving architectures such as hierarchies and networks.

How should constraints be released? It is too simplistic to design a constraint table and then follow the dictates of this structure. This amounts to simply

programming the observed behaviour of an infant directly. But there are much more subtle ways in which constraints may influence and modulate a learning system. We have experimented with emergent constraints and investigated simultaneous map generation and compared this with sequential generation and other schemes (Hülse and Lee 2010), and we find that explicit triggers for constraint lifting are not necessary (or even desirable). It is possible for constraints to be treated by these systems in an *emergent* way and then behavioural stages emerge as a consequence of the current state of the developing system. Further investigation is needed to understand this phenomena in terms of the literature surrounding developmental studies (Law et al. 2011).

How do the maps fit into a body-centric model of space? In our investigations, mappings have used individual reference frames. Those incorporating arm movements have been centred on the body centreline, whereas those incorporating visual saccades have been grounded in the visual space. To achieve composite tasks, such as visually guided reaching, we built mappings between these different reference frames. This opens up major questions about the structure of egocentric space and how an integrated sense of space can grow and be maintained. Such spaces must include tactile, visual and other representations and so become a kind of “ego-space”.

How can neural models of biological systems (e.g. the basal ganglia) be integrated into the LCAS framework? As our system is expanded, we expect situations to arise such as action selection, where a choice between alternative actions must be made. Maintaining the biological inspiration of the project, we see how neural models, such as those by Prescott et al. (2006) for action selection, can be incorporated. Further investigation is required to establish how other such models can be integrated into developmental approaches such as the LCAS architecture.

What is the role of novelty, and how does it relate to intrinsic motivation? We have used a single idea for intrinsic motivation: novelty, and a very elementary implementation of this. Research into novelty is not our main focus and this simple technique was used as a minimum complexity driver to provoke autonomous action. We appreciate that as competency increases more elaborate algorithms for novelty detection will be required. For example, our architecture cannot currently detect novelty in temporal events. We aim to integrate more sophisticated novelty detection algorithms such as those by Neto and Nehmzow (2007) and Oudeyer et al. (2007) in the future, but we also note that the whole question of intrinsic motivation is broader than novelty and may include other drivers.

10 Relation to Other Work

While there is now growing research activity in the area of developmental robotics, most related 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 throughout this chapter, particularly those that shed light on possible mechanisms and algorithms. But there is still a disproportionate 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 modelling of early development following a Piagetian approach has been that of [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 experiments by [Becker \(1973\)](#). Unfortunately, Drescher's implementation was a simulation with very primitive sensory-motor apparatus 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 and Maes 1994](#)).

A few models of infant grasping have been produced and some recent ones ([Oztop et al. 2004](#)) hint that visual guidance may not be central 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. \(2001\)](#) designed a neural architecture that computes motor commands from arm positions and desired directions. 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. 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 give identical performance, we wish to formulate general methods for constraint models and so favour more explicit algorithms.

11 Conclusions

The framework described in this chapter 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.

Support for our approach comes from various data. For example, studies of the order of cell activation in the foetus report that the first cells to be detected as active are the somatosensory cells, then the auditory, then visual, and finally, the multisensory cells become active ([Meredith et al. 1987](#)). This suggests what we have found that there are advantages if proprioception leads before vision 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 et al. 2004](#)). Reduction in degrees of freedom obtained by staged

development is also reported to be an effective strategy (Lungarella and Berthouze 2002; Sporns and Edelman 1993), as is the concept of constraints being beneficial to the emergence of stable patterns and helping to bootstrap later stages (Berthouze and 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 infant learning and adaptation can be very fast (Angulo-Kinzler et al. 2002; Rochat and Striano 1999) and in some cases only one trial or experience is needed to alter behaviour.

Our experiments have shown how stages in growth and behaviour may emerge from embodied agents through their exploration of the sensory-motor environment under constraint. It is important that behaviour grows and changes *without* any programming but through the shaping influence of experience on internal mechanisms. As an early researcher stated:

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

We must continue to search for mechanisms and supporting substrates that allow increasingly advanced behaviour to emerge from consolidated prior experience. Studies of constraint relationships are part of this endeavour and the role of constraints must be better understood.

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 longer term, we hope this will lead to new methodologies for building autonomous robot systems and better understanding and insights into human behaviour and growth.

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 experiential information. In R. C. Schank & K. M. Colby (Eds.), *Computer models of thought and language* (pp. 396–434). San Francisco: W.H. Freeman and Company.
- 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.

- Bodrova, E., & Leong, D. (2006). *Tools of the mind: the Vygotskian approach to early childhood education*. Columbus: Prentice Hall.
- Braitenberg, V., & Schüz, A. (1991). *Anatomy of the cortex: Statistics and geometry*. Berlin: Springer.
- Bruner, J. (1990). *Acts of meaning*. Cambridge: Harvard University Press.
- Bullock, D., & Grossberg, S. (1988). Neural dynamics of planned arm movements: emergent invariants and speed-accuracy properties during trajectory formation. *Psychology Review*, 95(1), 49–90.
- Caligiore, D., Ferrauto, T., Parsi, D., Accornero, N., Capozza, M., Baldassare, G. (2008). Using motor babbling and Hebb rules for modeling the development of reaching with obstacles and grasping. In *Proceedings of the international conference on cognitive systems (CogSys 2008)*, University of Karlsruhe, Karlsruhe, Germany, April 2–4, 2008.
- Campbell, R., & Bickhard, M. (1992). Types of constraints on development: an interactionist approach. *Developmental Review*, 12(3), 311–338.
- Carreira-Perpinan, M., Lister, R., Goodhill, G. (2005). A computational model for the development of multiple maps in primary visual cortex. *Cerebral Cortex*, 15(8), 1222–1233.
- Casey, B., Galvan, A., Hare, T. (2005). Changes in cerebral functional organization during cognitive development. *Current Opinion in Neurobiology*, 15(2), 239–244.
- Chao, F. (2009). *Constraint lifting and its application in developmental robotics*. PhD thesis, Department of Computer Science, Aberystwyth University, Wales.
- Chao, F., Lee, M., Lee, J. (2010). A developmental algorithm for ocular motor coordination. *Robotics and Autonomous Systems*, 58, 239–248.
- 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: MIT.
- Einarsdottir, H., Montani, F., Schultz, S. R. (2007). A mathematical model of receptive field reorganization following stroke. In *IEEE 6th international conference on development and learning* (pp. 211–216), Imperial College London, 11–13 July 2007, IEEE Computational Intelligence Society.
- Foner, L., & Maes, P. (1994). Paying attention to what's important: Using focus of attention to improve unsupervised learning. In *Proceedings of the 3rd international conference on simulation of adaptive behaviour* (pp. 256–265). Cambridge: MIT.
- Gallahue, D. (1982). *Understanding motor development in children*. New York: Wiley.
- Gasser, M., & Smith, L. B. (1998). Learning nouns and adjectives: a connectionist account. *Language and cognitive processes*, 13(2–3), 269–306.
- Gomez, G., Lungarella, M., Hotz, P. E., Matsushita, K., Pfeifer, R. (2004). Simulating development in a real robot: on the concurrent increase of sensory, motor, and neural complexity. In L. Berthouze, H. Kozima, C. G. Prince, G. Sandini, G. Stojanov, G. Metta, C. Balkenius (Eds.), *Proceedings of the fourth international workshop on epigenetic robotics: modeling cognitive development in robotic systems* (vol. 117, pp. 119–122). Lund University Cognitive Studies. Lund: LUCS.
- Goodhill, G., & Xu, J. (2005). The development of retinotectal maps: a review of models based on molecular gradients. *Network: Computation in Neural Systems*, 16(1), 5–34.
- Hainline, L. (1998). How the visual system develops. In A. Slater (Ed.), *Perceptual development: visual, auditory, and speech perception in infancy* (pp. 5–50). Hove: Psychology.
- Hendriks-Jensen, H. (1996). *Catching ourselves in the act*. Cambridge: MIT.
- Hülse, M., & Lee, M. H. (2010). Adaptation of coupled sensorimotor mappings: an investigation towards developmental learning of humanoids. In S. Doncieux, B. Girard, A. Guillot, J. Hallam, J.-A. Meyer, J.-B. Mouret (Eds.), *From animals to animats 11. Lecture notes in computer science* (vol. 6226, pp. 468–477). Berlin: Springer.
- Hülse, M., McBride, S., Law, J., Lee, M. (2010a). Integration of active vision and reaching from a developmental robotics perspective. *IEEE Transactions on Autonomous Mental Development*, 2(4), 355–367.

- Hülse, M., McBride, S., Lee, M. (2010b). Fast learning mapping schemes for robotic hand-eye coordination. *Cognitive Computation*, 2(1), 1–16.
- Hülse, M., McBride, S., Lee, M. H. (2009a). Implementing inhibition of return; embodied visual memory for robotic systems. In L. Cañamero, P.-Y. Oudeyer, C. Balkenius (Eds) *Proceedings of the ninth international conference on epigenetic robotics* (vol. 146, pp. 213–214). Lund University Cognitive Studies. ISBN 978-91-977-380-7-1
- Hülse, M., McBride, S., Lee, M. H. (2009b). Robotic hand-eye coordination without global reference: a biologically inspired learning scheme. In *IEEE 8th international conference on development and learning, (ICDL)* (pp. 1–6). New York: IEEE.
- Johnson, M. (1990). Cortical maturation and the development of visual attention in early infancy. *Journal of Cognitive Neuroscience*, 2(2), 81–95.
- Kalaska, J. (1995). Reaching movements: implications of connectionist models. In M. A. Arbib (Ed.), *The handbook of brain theory and neural networks* (pp. 788–793). Cambridge: MIT.
- Kalnins, I., & Bruner, J. (1973). The coordination of visual observation and instrumental behavior in early infancy. *Perception*, 2(3), 307–14.
- Kaplan, F., & Oudeyer, P.-Y. (2003). Motivational principles for visual know-how development. In C. G. Prince, L. Berthouze, H. Kozima, D. Bullock, G. Stojanov, C. Balkenius (Eds.), *Proceedings of the third international workshop on epigenetic robotics* (vol. 101, pp. 73–80). Lund University Cognitive Studies. ISBN 91-974741-X.
- Kaplan, F., & Oudeyer, P.-Y. (2007). In search of the neural circuits of intrinsic motivation. *Frontiers in Neuroscience*, 1(1), 225–236.
- Keil, F. (1990). Constraints on constraints: Surveying the epigenetic landscape. *Cognitive Science*, 14(4), 135–168.
- Law, J., Lee, M. H., Hülse, M., Tomassetti, A. (2011). The infant development timeline and its application to robot shaping. *Adaptive Behaviour*, 19(5), 335–358.
- 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 F. Kaplan, P.-Y. Oudeyer, A. Revel, P. Gaussier, J. Nadel, L. Berthouze, H. Kozima, C.G. Prince, C. Balkenius (Eds), *Proceedings of the sixth international workshop on epigenetic robotics* (vol. 128, pp. 55–62), September 20–22, 2006, Paris, France, Lund University Cognitive Studies. ISBN 91-974741-6-9
- Lee, M. H., Meng, Q., Chao, F. (2007a). Developmental learning for autonomous robots. *Robotics and Autonomous Systems*, 55(9), 750–759.
- Lee, M. H., Meng, Q., Chao, F. (2007b). Staged competence learning in developmental robotics. *Adaptive Behaviour*, 15(3), 241–255.
- Lungarella, M., & Berthouze, L. (2002). Adaptivity through physical immaturity. In C.G. Prince, Y. Demiris, Y. Marom, H. Kozima, C. Balkenius (Eds) *Proceedings of the second international workshop on epigenetic robotics* (vol. 94, pp. 79–86). Lund University cognitive studies, August 10–11, 2002, Edinburgh, Scotland. *Studies*, 94. Lund: LUCS
- Lungarella, M., Metta, G., Pfeifer, R., Sandini, G. (2003). Developmental robotics: a survey. *Connection Science*, 15(4), 151–190.
- Mallot, H., Von Seelen, W., Giannakopoulos, F. (1990). Neural mapping and space-variant image processing. *Neural Networks*, 3(3), 245–263.
- Martinetz, T. (1993). Competitive Hebbian learning rule forms perfectly topology preserving maps. In *Proceedings of the ICANN'93, International conference on artificial neural networks* (vol. 93, pp. 427–434). Springer.
- Meng, Q., & Lee, M. H. (2007). Automated cross-modal mapping in robotic eye/hand systems using plastic radial basis function networks. *Connection Science*, 19(1), 25–52.
- 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.
- Metta, G., Craighero, L., Fadiga, L., Ijspeert, A., Rosander, K., Sandini, G., Vernon, D., von Hofsten, C. (2009). A roadmap for the development of cognitive capabilities in humanoid robots. Technical Report D2.1, University of Genoa.

- Nakayama, K., & Silverman, G. (1986). Serial and parallel processing of visual feature conjunctions. *Nature*, 320(6059), 264–265.
- Neto, V., & Nehmzow, U. (2007). Visual novelty detection with automatic scale selection. *Robotics and Autonomous Systems*, 55(9), 693–701.
- Oudeyer, P., Kaplan, F., Hafner, V. (2007). Intrinsic motivation systems for autonomous mental development. *IEEE Transactions on Evolutionary Computation*, 11(2), 265–286.
- 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.
- Pfeiffer, R., & Bongard, J. (2006). *How the body shapes the way we think*. Cambridge: MIT.
- Piaget, J. (1973). *The child's conception of the world*. London: Paladin.
- Piek, J. P. (2002). The role of variability in early motor development. *Infant Behavior and Development*, 25(4), 452–465.
- Piek, J. P., & Carman, R. (1994). Developmental profiles of spontaneous movements in infants. *Early Human Development*, 39(2), 109–126.
- Pouget, A., & Snyder, L. (2000). Computational approaches to sensorimotor transformations. *Nature Neuroscience*, 3, 1192–1198.
- Prescott, T., Montes González, F., Gurney, K., Humphries, M., Redgrave, P. (2006). A robot model of the basal ganglia: behavior and intrinsic processing. *Neural Networks*, 19(1), 31–61.
- Prince, C., Helder, N., Hollich, G. (2005). Ongoing emergence: a core concept in epigenetic robotics. In L. Berthouze, F. Kaplan, H. Kozima, H. Yano, J. Konczak, G. Metta, J. Nadel, G. Sandini, G. Stojanov, C. Balkenius (Eds.) *Proceedings of the fifth international workshop on epigenetic robotics* (vol. 123, pp. 63–70). Lund University cognitive studies. ISBN 91-974741-4-2.
- Redding, G., & Wallace, B. (2006). Generalization of prism adaptation. *Journal of Experimental Psychology: Human Perception and Performance*, 32(4), 1006–1022.
- 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.
- Schmidhuber, J. (1990). Learning algorithms for networks with internal and external feedback. In *Proceedings of the 1990 connectionist models summer school* (pp. 52–61), San Mateo, CA.
- Smith, L., & Gasser, M. (2005). The development of embodied cognition: six lessons from babies. *Artificial Life*, 11(1–2), 13–29.
- Spelke, E. (1998). Nativism, empiricism, and the origins of knowledge. *Infant Behavior and Development*, 21(2), 181–200.
- 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.
- Sterling, P. (1999). Deciphering the retina's wiring diagram. *Nature Neuroscience*, 2, 851–852.
- Thelen, E. (1995). Motor development. *American Psychologist*, 50, 79–95.
- 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, pp. 243–277). New York: Lawrence Erlbaum Associates, Inc.
- 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.

- Ververs, I., Gelder-Hasker, V., Marja, R., De Vries, J., Hopkins, B., Van Geijn, H. (1998). Prenatal development of arm posture. *Early Human Development*, *51*(1), 61–70.
- Westermann, G., & Mareschal, D. (2004). From parts to wholes: mechanisms of development in infant visual object processing. *Infancy*, *5*(2), 131–151.
- Wurtz, R., & Goldberg, M. (1972). The primate superior colliculus and the shift of visual attention. *Investigative Ophthalmology & Visual Science*, *11*(6), 441–450.