970

Ontogenetic Constraints on Scaling-Up Sensory-Motor Systems

Julie C. Rutkowska

School of Cognitive & Computing Sciences
University of Sussex
Brighton
BN1 9QH, U.K.

julier@cogs.susx.ac.uk

Abstract

Significant new work in AI sees perceptual and motor skills as the 'hard' problems solved by intelligent systems, solutions to which impose important constraints on remaining components of natural intelligence. This paper outlines some ways in which an ontogenetic focus may prove relevant to attempts to understand adaptive behaviour in mobot/robot-environment systems. Synchronic accounts that seem to work for basic abilities may fare less well when required to accommodate diachronic processes of change; and understanding the kinds of global changes that routinely occur as a consequence of natural sensory-motor systems' local interactions with the environment can contribute to bridging the explanatory gap between such situated systems and their more abstract conceptual counterparts. Three main issues are discussed that focus on if and how scaling-up from basic to supposedly higher abilities is possible: (1) Whether early infant abilities are better explained in terms of conceptual representations or a computational account of action; (2) Reciprocal constraints between cognitive and physical-motor mechanisms, the role they play in paradigmatic cases of adaptive change, and the need to take seriously the internal structure of behaviour; (3) How far typical self-organizing connectionist networks take us towards understanding a system that is capable of mapping recurrent viable patterns of activity into more permanent adaptive changes.

971

The scaling-up problem

Significant new work in AI sees perceptual and motor skills as the 'hard' problems solved by intelligent systems, solutions to which will impose serious constraints on remaining components of natural intelligence (e.g. Brooks, 1991; Meyer & Guillot, 1990); and the developmental psychology of infancy is increasingly taking action seriously in its own right as the foundation of an autonomous pragmatic knowledge (e.g. Hobson, 1991; Mounoud, 1982; Rutkowska, 1991). This paper outlines some ways in which the ontogenetic focus of this direction in developmental psychology may prove relevant to AI's recent attempts to understand adaptive behaviour in mobot/robot-environment systems.

The shared assumption that intelligent abilities are grounded in a system's sensory-motor exchanges with its environment raises the recurrent problem of whether and how scaling-up from basic to supposedly higher abilities is possible. To what extent can recent AI work do away altogether with classical symbolic notions of computation? And with regard to what kind(s) of mechanisms might it be necessary to invoke concepts and so forth? Bridging the gap is made problematic by a dearth of paradigmatic examples around which discussion may be framed. Analysis drawing on implemented AI examples tends to juxtapose the functioning of situated systems capable of, say, task-specific behaviours like moving towards a target or avoiding bumping into obstacles (increasingly thought of as concept- and representation-free) with that of systems exhibiting complex sequential activities that are hard to decouple from a traditional centralized, disembodied view of general-purpose reasoning (hence of conceptualizing the world and manipulating exhaustive central representations of it to preplan activity). Rapprochement between such extremes seems unlikely. However, the developmental psychology of human infancy offers relevant and instructive examples that feature changing organization in everyday activities. The kinds of adaptive change that infants routinely exhibit in activities such as prehension far outstrip the scaling-up potential of current AI sensory-motor systems, and they help to clarify the principles that psychologically valid cases of such systems will need to obey.

Approaching epigenesis through conceptual networks or action

Mainstream (Anglo-Saxon) developmental psychology currently attributes a key place to 'epigenesis' (e.g. Carey & Gelman, 1991). Despite this, prevailing theories are neither *enactive* (Varela, 1988) nor *epigenetic* (Piaget, 1953; 1971). Central to the latter (European-influenced) perspectives is a systematic style of constructivism, where the system concerned spans the subject's (perceptual and behavioural) activity in the environment. Psychological structures that are enacted ('brought forth') or develop as a consequence of that activity are an outcome of mutually generated organization; they originate neither in the subject nor in the environment, and their construction cannot be meaningfully decomposed into contributions from the subject (endorsing nativism) and/or an objective environment (endorsing empiricism).

In contrast with this, infants' understanding of objects is increasingly being attributed to a precocious object concept, with a central conceptual system responsible for parsing an unsegmented perceptual array into discrete objects on the basis of unchanging principles such as cohesion, boundedness and rigidity (Spelke, 1990, 1991). This fits a general style of developmental explanation in which domain knowledge is attributed to an explicitly represented network of concepts, whose origins tend to be located well towards the nativist end of what is essentially a nativist-empiricist scale, in line with the reasoning of Chomsky (1980) and Fodor (1983). Development in such a system is viewed as an abductive mapping from environments to mental representations (Keil 1981), but this way of viewing development 'solves' the scaling-up problem simply by sleight of hand. It establishes an internal-external dichotomy that (mis)localizes key constraints on adaptive behaviour -- in either internal cognitive structures or external task-domain ones. Activity such as overt behaviour is at best treated as a convenient index of central knowledge.

Grounding infant abilities in action challenges the appropriateness of attributing concepts to the infant ab initio, and suggests that to do so presents a molar, misleading view of underlying mechanisms and of the kinds of change of which they are capable. A typical definition of strictly conceptual functioning is provided by Kirsh's (1991, p.163) focus on *predicability:* 'to have a concept is, among other things, to have a capacity to find an invariance across a range of contexts, and to reify that invariance so that it can be combined with other appropriate invariances.....to identify the common property which two or more objects share and to entertain the possibility that other objects also possess that property.' Neither aspect of such definitions -- explicit representation of invariances as properties of something(s); and flexibility and interrelatedness of knowledge associated with syntactic combinatorial capacities of the kind defined by the generality constraint (Evans, 1982) -- may offer a valid characterisation of initial infant mechanisms. The global functional organization of those mechanisms can, however, be better understood in classical computational terms that subsequently prove useful for phrasing developmental issues (Rutkowska, 1991, 1992).

If we take the example of early prehension, the conditions under which infants reach for things, and their typical errors, strongly implicate reliance on a visually provided (digital-symbolic) description in a representational format akin to Marr's (1982) 2½D sketch, which purports to make explicit the depth and orientation of patches of visible surface and their discontinuities. For example, infants reach not for the smallest or closest of two objects, but for segmented surfaces whose boundaries are distinguished from others by separation in depth and/or motion (von Hofsten & Spelke, 1985). If these conditions are violated, say by placing a reached-for object directly onto a larger one, reaching is disrupted (Wishart & Bower, 1984). A range of infant abilities can be related to other forms of description (cf. Rutkowska, 1991, in press). Marr fails adequately to consider the role of the behavioural component of action in perceptual processing, but his general model does suggest a useful basis for preattentive vision. More 'central' computational mechanisms that are additionally required to account for performance do not need to operate on the descriptions provided by low-level vision to generate integrated, hierarchical descriptions of objects and their properties. Instead, they can be viewed more pragmatically in terms of action programs: virtual mechanisms whose operation selectively exploits task-relevant aspects of multiple descriptions (via processes such as pattern matching and variable binding) to support the direct invocation of behavioural procedures.

If infants' object knowledge is grounded, along these lines, in low-level visual processing, it need not involve concepts in the technical sense defined above. Neither the descriptions that this kind of processing constructs, nor the processes that exploit them, are syntactic in the necessary sense. Making explicit an aspect of the physical world, such as a surface, over many situations does not entail any ability to represent it as a property that is common to that range of situations, let alone potentially applicable to others. Visual processing does not deliver a description of anything as, or with the property of being, a surface; it simply delivers or fails to deliver a description of a surface. The visually based knowledge of surface segmentation that supports infants' organized behaviours towards objects is not available in a form that would allow a central system explicitly to represent properties of an object, e.g. that 'it' is bounded or cohesive, let alone to entertain the idea that these might be properties of other, arbitrary things.

The organization of infant action mechanisms nevertheless fits two important ideas relevant to understanding action-based representation as a means of establishing selective correspondences with the world, rather than as a model that substitutes for it. One is the notion of *attunement to constraints* that is developed by Israel in the situation-semantics framework. This proposes that 'one situation (that is what is going on at one spatiotemporal location) carries information about another situation relative to a law or constraint relating types of situations (1988, p.124)', but a subject can *carry* information without *having* that information. For example, an intra-subject situation such as a description of red spots produced by the visual system can be said to carry the information that a child has measles; but the subject cannot be said to have that information in the absence of information and control states 'that both carry the information about the spots and also lead to behaviour approriate to the tale they are telling about the infection -- appropriate, that is, given the goal of a healthy, happy child (1988, p.124).' Perceptual-behavioural integration, which involves perceptually provided information being not only available to, but also explicitly exploited by, the action program level, is thus significant to the subject's understanding of, and ability to satisfy, constraints on the way the world works.

Additionally, the organization of the key computational components of this kind of action-based representation fit the criteria for a *natural system of representation* in Dretske's sense. What a system represents is not determined by what its expressive elements (in the present example, descriptions constructed by low-level vision) may indicate or mean in isolation; it is determined by what those elements have the function of indicating or meaning. Unconventional natural systems of representation are distinguished from its conventional counterparts (e.g. maps or musical notation) in that they 'have *their own* intrinsic indicator functions, functions that derive from the way the indicators are developed and used *by the system of which they are part* .....These functions are not assigned. They do not depend on the way others may use or regard the indicator elements (Dretske, 1988, p.62).' For instance, the way they are selectively exploited by action programs to generate behaviour allows similar low-level visual elements associated with directional selectivity to have the function of indicating an approaching object on a hit course for the human infant (Ball & Tronick, 1971), food for the frog (Lettvin, Maturana, McCulloch & Pitts, 1958), and a potential mate for the hoverfly (Collett & Land, 1978).

Serendipity: from viable activity to intentional action

The approach to infant ability that was sketched in the preceding section sees it as heavily dependent on the way that subject and environment operate together to generate and sustain viable patterns of activity. One advantage of invoking the classical computational virtual-physical machine distinction to conceptualize behaviour is that it offers a descriptive language for beginning to explore the key issue of reciprocal constraints between cognitive and physical-motor structures, and of the role they play in the kinds of adaptive change that we call 'development'. The notion of a virtual machine that can implement a computational process in a physical machine extends beyond program-mind and computer-brain analogies, since the processes of the central nervous system are not simply coextensive with what goes on in the brain. The body is equally relevant, and viewing behaviour in terms of program procedures that govern processes including movements of eyes, hands, and so forth offers a vital way of conceptualizing the causal, spatio-temporal aspect of its functioning in the world (besides the teleological aspect, to which traditional accounts pay most heed).

To the extent that reasoning is involved in such a system, it is more akin to *situated inference* (Barwise, 1987) than to *formal inference* of the style traditionally associated with computational accounts (Fodor, 1980). The validity of situated inference relies on the embedding circumstances of the subject, not simply on the execution of appropriate central rules. Hence, the soundness of infants' typical inferences (e.g. when it is appropriate to generate avoidance behaviour) depends on the reliability of their action-based representation, which depends on the continuation of natural environmental conditions. If those conditions are altered or break down, such inferences will no longer be valid, even if identical computational steps have been executed (e.g. faced with unnatural conditions such as a psychologist's laboratory, an infant will inappropriately attempt to avoid an expanding shadow).

Exploiting constant environmental features is a basic kind of situated inference. In general, the operating principles of embodied naturally intelligent systems involve: embedding; externalization of process (via the dual nature of overt behaviour); transaction; and scaffolding by physical and social environments. The implications of this for both synchronic and diachronic functioning is instructive about adaptive behaviour. In its synchronic sense, adaptive behaviour is defined primarily in terms of its ability to handle changing and often unpredictable environmental circumstances. But it can also be considered in a diachronic sense, where its defining characteristic looks predominantly like the construction of predictability in order to exploit environmental structure for anticipatory and planful activity, hence economising on memory and processing resources. The restructuring of action mechanisms that underlies such developmental change supports performance that begins to look as if it requires a 'conceptual' level of functioning.

The notion of *serendipity* -- accidental, fortunate discovery -- offers a way of starting to think about the role of behavioural processes in such adaptive change. It involves the subject not only getting things done with environmental support that they could not achieve alone (the key function of scaffolding) but getting things done that they did not 'know' how to do and were not even attempting to get done.

For example, Winograd's (1987) hypothetical pizza-delivery robot, Ramona, offers a clear (synchronic) example of how physical properties of the subject can constrain activity by meshing with those of the embedding environment to undercut the need for symbol structures that provide even an implicit action-based task representation. Pizza orders take the form of letters making up words in a standard vocabulary, and Ramona enters them into horizontal rows in a grid in its computational field. The lowest row dropping into an output buffer indicates which pizza should be prepared next; only when all are ready can the order be delivered. Since bigger pizzas take longer to cook, preparation sequence should obviously be taken into account, and Ramona gives the illusion of having this under control. But Ramona's behaviour is not, in fact, designed for this. It 'comes for free' as a side-effect of what is, from the robot's perspective, the genuinely task-achieving behaviour of putting names into its grid.

The most relevant behavioural process under the control of Ramona's action program is just something like 'raise your arm until it gets to the grid'; how high up it reaches is determined by exchanges between extra-program physical properties of subject and environment: bigger pizzas have longer names, so involve a greater hence heavier load of letters; and the robot's arm is floppy, the height to which it rises at the grid determined by the weight it carries. Hence, heavier names end up nearer the bottom of the grid and the preparation of bigger pizzas starts before that of smaller ones.

A key issue for an epigenetic account of development is concerned is whether the patterns of viable activity established in such serendipitously successful movements could come under action program control to attain the status of purposive, task-oriented behaviours? In Ramona's case, it is unlikely. Even if its sensory capacities enabled it to detect the contingency, there is no range of variation in arm rigidity that could be controlled in this way to enable it to raise its arm in particular ways in order to affect sequencing. In the case of the human infant, however, variation along physical dimensions such as force, amplitude and velocity are available, and such anticipation routinely emerges. A good example is provided by the way grasping and lifting objects is restructured to exhibit permanent, foresightful adaptation to covariation between their size and weight (Mounoud & Hauert, 1982a & b).

6- to 8-month-old infants are capable of effective grasping, whatever the object's specific weight, through a local context-specific adaptation to particular circumstances. This type of grasp is just as effective if an object that has been grasped several times is replaced by a much lighter one of identical size. By way of contrast, the grasping of older infants shows disruption, including sudden upward arm movement, in the face of such trick substitutions. Subjects shake and closely inspect the objects concerned, as if aware of the anomaly inherent in substitution trials; i.e. that some examples depart from a 'norm' or 'general rule'. By 14-16 months, there is disruption in the trick context, but it is compensated for by reversion to the original type of organization. These infants appear to have established a pragmatic understanding of size-weight covariation. Reorganization following on the outcomes of perceptual and behavioural functioning in a range of grasping situations underlies a new ability to 'infer' an object's weight from its perceived size, and to contract the hand appropriately, in an anticipatory way, prior to grasping.

The same 3 levels of organization are found with many ages, domains of knowledge and types of representation. The general nature of the process revealed is one of discovering and making explicit recurrent patterns that arise from viable activity in the environment. It can be thought of as *attunement* to novel constraints on successful action; or as *endogenously motivated assignment* of a novel indicator function to certain dimensions of visually produced descriptions. Relevant to the idea of what conceptual functioning might entail, such examples reveal increasing systematicity of the subject's knowledge, to the extent that action-based representations appear to make explicit the general conditions involved in initially purely practical success across a range of individual contexts. Appreciation of anomaly suggests a regularity across a range of situations comes to acquire the status of an expected property of those and similar situations.

A key implication of such phenomena is that explorations of intelligence along mobot/robot lines need to take behaviour even more seriously. In this example of action-based representation, change involves the way future grasping behaviours will be programmed, and the infant's new knowledge of weight does not lie in perception or in behaviour -- it is distributed through the system. The new behavioural form 'selectively models' aspects of the environment no less than perception might. Descriptions produced by the perceptual component of action may indicate something about the environment, due to their correspondence(s) with it, and to the role they play within the system. But there is a sense in which behavioural procedures (the symbol structures that govern motor processes) do so too -- through the way their function is to generate processes that establish correspondence with spatio-temporal aspects of the environment.

In contrast with this emphasis, even computational approaches that endorse causal embedding tend to focus almost exclusively on the perceptual component of action. For example, Harnad's (1992) 'hybrid' solution to the problem of *symbol grounding* argues that the right place to look for how a machine's symbols can mean something to it, and not just to the person who designed it, is in its robotic capacity. In particular, the transducer/effector and other analogue structures and processes that constitute its interface with its environment are not just 'peripherals' whose implementation is no more than a matter of convenience: 'you have to *be* the transducer/effector and analogue structures and processes (not just the symbolic ones) in order to be a mind: It is not that the mind *receives* the transducer/effector or analogue activity (or, for that matter, the symbolic activity) as data. If the mind is grounded this way then it just *is* the activity of those structures and processes (Harnad, 1992, p.80).' The details of Harnad's analysis focus exclusively on how connectionist computation could support a sensory-category learning mechanism that is capable of producing categorical representations of the kind implicated in sorting and labelling objects. But the way both sensory-transducers and effectors mesh with the environment determines what robotic activities 'work', hence the effector side of action ought not to be ignored in the kind of supervised learning that Harnad considers essential.

Some AI accounts are beginning to take behaviour more seriously at a synchronic level. For example, in the nascent *animate vision* paradigm, behaviour like gaze control (cf. Piaget's (1953) 'looking' or 'directing the glance') is being conceptualized as playing a vital role in vision algorithms (Ballard, 1989). However, fleshing out a developmental perspective is likely to need a better computational framework for exploring the internal structure of behaviour. Even

Ballard's (1993) behaviourally sophisticated extension of animate vision to the modelling of eye-hand coordination locates almost all of its internal complexity with details of sensory inputs, how fixation behaviour affects their representation, and their selective use by attentional processes. Behaviours themselves appear to be treated as molar entities that lack internal structure, e.g. 'grasp' or 'fixate'. We currently lack a behavioural-motor equivalent of the computational framework available for exploring the organization of 'autonomous computational reflexes' of low-level, preattentive visual processing. It is plausible, however, that there is a corresponding level of analysis at which preadapted movement mechanisms have their own computational principles of operation, distinct from the action program level.

The best place for connectionist networks

A classical computational way of thinking about the kinds of developments sketched above is in terms of the construction, at the action program level, of new patterns including variables that derive their values from preattentive descriptions of object dimensions and are implicated in the control of behavioural procedures governing the movements involved in action. In a very traditional system such as Sussman's (1975) blocksworld learning-program, HACKER, for example, a process of *variablization* underlies its ability to make explicit general patterns in its (sensory and motor) activity that enable it to generalize available solutions to new problems. Such ideas may seem outmoded and too high-level to be of much interest. There is increasing enthusiasm for purportedly *subsymbolic* connectionist styles of computation, as total alternatives (e.g. Churchland, 1986, 1989) or as more useful microtheories for which classical phenomena may offer a macrotheory (e.g. Smolensky, 1987). From a conciliatory perspective, symbolic phenomena are seen as an emergent property of distributed connectionist systems, but the latter hold out the prospect of a clearer understanding of the relation between computational and physical accounts of intelligent functioning.

How much farther do typical self-organizing connectionist systems take us towards understanding the workings of a system that is capable of mapping recurrent viable patterns of activity into more permanent adaptive changes to its processing potential? It seems likely that connectionist networks may provide the implementation for components of a system whose global functional organization can be better understood in terms of classical computational structures and processes.

Self-organizing connectionist networks seem attractive to developmental biologists because of the way that global structure can be shown to result from the operation of local rules defined in terms of the activation values of, and connection strengths between, their multiple simple units (e.g. Stewart, 1989). They have also been endorsed from the viewpoint of developmental psychology, in an attempt to devise a genuinely epigenetic account of knowledge acquisition. This is based on attempts to conceptualize an endogenous process of *representational redescription* (Karmiloff-Smith, 1991), in which representations underlying knowledge of novel domains are constructed not from environmental inputs but from representations responsible for subjects' successful practical activity.

Clark and Karmiloff-Smith (in press) offer a computational account of how knowledge that is initially implicit in successfully functioning activity could become explicit, in terms of a distinction between first- and second-order connectionist systems. As a typical example of a first order system, they consider Sejnowski and Rosenberg's NETtalk, which takes text as input and outputs speech. This is characterized by exhibiting example-driven learning, through which it becomes trained to statistical regularities between inputs and outputs. The only rules it comes to know are implicit, emerging from interactions between its units. We might describe some of NETtalk's successes in text to speech transformation as coming to 'know about vowels', for example, but this abstraction concerning what its hidden units are doing is nowhere available to the system itself. Such systems, are thought of as a suitable basis for the acquisition of example-driven, implict knowledge. However, unless the problem domain stays stable, they are limited by being unable to change strategy without retraining, which would destroy the organization they have built-up.

Clark and Karmiloff-Smith suggest that flexibility, i.e. systematically manipulable components that support example-independent generalizations, could be achieved via a second-order connectionist system. Far fewer examples of candidate systems exist for this category, but a possible contender is Mozer and Smolensky's (1989) *skeletonization* technique. In contrast with first-order networks, which change their organization through additional training with further inputs, hence losing what learning has already occurred, this is able to use a procedure that operates on a (first-order) network that has successfully been trained already. The procedure identifies the input and hidden units that are most important in determining performance, keeps those and deletes those that are least relevant. The 'skeletonized' network proves to be better at generalizing, learns faster and its pared-down structure can be described 'in terms of a small number of rules instead of an enormous number of parameters (Mozer and Smolensky, 1989, p.15).'

Two main observations about these ideas are pertinent here. The first queries the extent to which unstructured first-order connectionist systems can be unequivocally endorsed as a basis for the acquisition of initially implicit knowledge in successful activity. Like Harnad's (1992) example of a system learning about the property of line-length, the majority of connectionist examples, which impress due to features like graceful degradation, are primarily attempting 'objective' categorization that does not consider how to achieve a system capable of flexible parsing and exploitation of (preattentive) perceptually generated descriptions with regard to their relevance to activity. (And the selectivity of trainers' input examples casts doubt on the significance of their more limited achievement (van Brakel, 1991)).

Once abilities involving coordinations of multiple procedures are at issue, it seems highly probable that more structure must be imposed on the 'untrained' system and its exchanges with the environment. Even early infant acquisitions are of this form, and it is clearly exhibited in means-end coordinations, such as learning to remove a cover in order to search for a hidden object that is initially found by serendipitous means. Norris' (1991) example of a date-calculation task that involves a 3-step algorithm found it impossible to train an unstructured network except on each individual step, one at a time. Effectively, its eventual learning and generalization was made possible only by being told about the structure of the task and the algorithm.

It is also notable that what skeletonization is doing in a connectionist network looks very similar to classical machine-learning processes like variablization. Even though an anti-variable, anti-pattern matching (and even anti-neural network) stance is often seen as typical of recent mobotic research such as that of Brooks (1991), it is interesting that his *subsumption architecture* for layered control of independent task-achieving behaviours bears many similarities to the sketch of an action-program governed system sketched here, despite Brooks being at pains to talk of wires and engineering considerations. For instance, a set of action programs that can exploit the same visual descriptions and motor processes fits with Brooks' rejection of the traditional clear distinction between serially ordered perception, central and action subsystems, particularly insofar as there is no homogeneous central representation of the world. Perceptual processing that constructs multiple descriptions of the environment, each accessible to different action programs, rather than an integrated object description, fits Brooks' claim that there is no single (central) place where a description of the world is delivered. It is also compatible with his assumption that a number of algorithms from different layers may be run on the same sensor outputs (cf. Brooks' notion of higher layers exploiting the motor processes of lower ones.)

In a similar vein, there are interesting parallels between such action program ideas and the functional decomposition of Ballard's (1993) model of animate vision in action. Ballard talks of this as a subsymbolic system, but is using this term in an unusual sense, broadly to imply computation that takes account of physical properties such as body form and movement. But considering computational processes to include body movements is, as noted earlier, a natural extension of the classical virtual-physical machine distinction, which considers symbol structure manipulation as a key component of behavioural control. This idea is orthogonal to the semantically central notion of a subsymbolic level relying on dimension-shifted representations (Clark & Lutz, 1992), athough, to the extent that a constraint analysis/relaxation style of computation is implicated in both visual processing and movement generation, it is not incompatible with it. The multiple simple units of connectionist connectionist networks may be viewed more appropriately as another virtual machine level than as a direct line to the subject's physical structure and causal embedding in the physical environment.

Such considerations do not deny that the kinds of global ontogenetic change that have been sketched in this paper probably depend upon more local exchanges between the structure (at a number of levels) of subject and environment. But any idea that connectionist thinking offers a ready-made route to explaining the developmental interplay between physical-motor and cognitive levels of structure should currently be rated overoptimistic.

References

Ball, W. and Tronick, E. (1971) Infant responses to impending collision: Optical and real. *Science* **171**, 818-820.

Ballard, D.H. (1989) Reference frames for animate vision. *Proceedings of the International Joint Conference on Artificial Intelligence.* Volume 2. Pp. 1635-1641.

Ballard, D.H. (1993) Sub-symbolic modelling of Hand-eye coordination. In D.E. Broadbent (ed.) *The Simulation of Human Intelligence*. Oxford: Blackwell.

Barwise, J. (1987) Unburdening the language of thought. *Mind and Language* **2**, 82-96.

Brooks, R.A. (1991) Intelligence without representation. *Artificial Intelligence* **47**, 139-160

Chomsky, N. (1980) Rules and representations. *Behavioral and Brain Sciences* **3**, 1-61.

Churchland, P.M. (1986) The continuity of philosophy and the sciences. *Mind and Language* **1**, 5-15.

Churchland, P.M. (1989) On the nature of explanation: A PDP approach. In P.M. Churchand *The Neurocomputational Perspective.* Cambridge, Mass.: M.I.T. Press.

Clark, A. and Lutz, R. (eds.) (1992) *Connectionism in Context.* London: Springer-Verlag.

Clark, A. and Karmiloff-Smith, A. (in press) The cognizer's innards: A psychological and philosophical perspective on the development of thought. *Mind & Language.*

Collett, T.S. & Land, M.F. (1978) How hoverflies compute interception courses. *Journal of Comparative Physiology A* **125**, 191-204.

Dretske F.I. (1988) *Explaining Behaviour: Reasons in a World of Causes.* Cambridge, Mass.: Bradford/M.I.T. Press.

Fodor, J.A. (1980) Methodological solipsism considered as a research strategy in cognitive science. *Behavioral and Brain Sciences* **3**, 417-214.

Fodor, J.A. (1983) *The Modularity of Mind: An Essay on Faculty Psychology.* Cambridge, Mass.: Bradford/M.I.T. Press.

Harnad, S. (1992) Connecting object to symbol in modelling cognition. In Clark, A. and Lutz, R. (eds.) *Connectionism in Context.* London: Springer-Verlag.

Hobson, P. (1991) Against the theory of 'Theory of Mind'. *British Journal of Developmental Psychology* **9**, 33-53.

Israel, D. (1988) Bogdan on information. *Mind and Language* **3**, 123-140.

Karmiloff-Smith, A. (1991) Beyond modularity: Innate constraints and developmental change. In S. Carey and R. Gelman (eds.) *The Epigenesis of Mind: Essays on Biology and Cognition*. Hillsdale, N.J.: Lawrence Erlbaum.

Keil, F.C. (1981) Constraints on knowledge and cognitive development. *Psychological Review* **88**, 197-227.

Kirsh, D. (1991) Today the earwig, tomorrow man? *Artificial Intelligence* **47**, 161-184.

Lettvin, J., Maturana, H., McCulloch, W. and Pitts, W.S. (1958) What the frog's eye tells the frog's brain. *Proceedings IRE* **147**, 1940-1951.

Marr, D. (1982) *Vision*. San Francisco: Freeman.

Meyer, J.-A. and Guillot, A. (1990) Simulation of adaptive behaviour in animats: Review and prospect. In J.-A. Meyer and S.W. Wilson (eds.) *From Animals to Animats: Proceedings of the First International Conference on the Simulation of Adaptive Behaviour*.

Mounoud, P. (1982) Revolutionary periods in early development. In T.G. Bever (ed.) *Regressions in Mental Development: Basic Phenomena and Theories*. Hillsdale, N.J.: Lawrence Erlbaum.

Mounoud, P. and Hauert, C.A. (1982a) Development of sensorimotor organization in young children: Grasping and lifting objects. In G.E. Forman (ed.) *Action and Thought: From Sensorimotor Schemes to Thought Operations*. New York: Academic Press.

Mounoud, P. and Hauert, C.A. (1982b) Sensorimotor and postural behaviours: Their relation to cognitive development. In W.W. Hartup (ed.) *Review of Child Development Research*. Volume 6. Chicago and London: U Chicago Press.

Mozer, M and Smolensky, P. (1989) Using relevance to reduce network size automatically. *Connection Science* **1**, 3-17.

Norris, D. (1991) The constraints on connectionism. *The Psychologist* **4**, 293-296.

Piaget, J. (1953) *The Origin of Intelligence in the Child*. London: Routledge and Kegan Paul.

Piaget, J. (1971) *Biology and Knowledge*. Edinburgh University Press.

Rutkowska, J.C. (1990) Action, connectionism and enaction: a developmental perspective. *AI & Society* **4**, 96-114.

Rutkowska, J.C. (1991) Looking for 'constraints' in infants' perceptual-cognitive development. *Mind and Language* **6**, 215-238.

Rutkowska, J.C. (in press) *The Computational Infant: Looking for Developmental Cognitive Science*. Hemel Hempstead: Harvester Wheatsheaf.

Smolensky, P. (1987) Connectionist AI, symbolic AI and the brain. *Artificial Intelligence Review* **1**, 95-110.

Spelke, E.S. (1990) Principles of object perception. *Cognitive Science* **14**, 29-56.

Spelke, E.S. (1991) Physical knowledge in infancy: Reflections on Piaget's theory. In S. Carey and R. Gelman (eds.) *The Epigenesis of Mind: Essays on Biology and Cognition*. Hillsdale, N.J.: Lawrence Erlbaum.

Stewart, J. (1989) A biologist's view of connectionism. Paper presented at the College International de Philosophie and Universite de Technologie de Compiegne Seminar, 'La Cognition'. Chantilly, France, July.

Van Brakel, J. (1991) Meaning, prototypes and the future of cognitive science. *Minds and Machines* **1**, 233-257.

Varela, F.J. (1988) *Cognitive Science: A Cartography of Current Ideas*. Unpublished translation of F.J. Varela (1989) *Connaitre -- Les Sciences Cognitives: Tendances et Perspectives*. Paris: Editions du Seuil.

von Hofsten, C. and Spelke, E.S. (1985) Object perception and object-directed reaching in infancy. *Journal of Experimental Child Psychology* **114**, 192-212.

Winograd, T. (1987) Cognition, attunement and modularity. *Mind and Language* **2**, 97-103.

Wishart, J.G. and Bower, T.G.R. (1984) Spatial relations and the object concept: A normative study. In L.P. Lipsitt and C. Rovee-Collier (eds.) *Advances in Infancy Research*. Vol. 3. Norwood, N.J.: Ablex.