Article

What's Special About the Development of the Human Mind/Brain?

ANNETTE KARMILOFF-SMITH AND ANDY CLARK

There's always a certain frustration (articles can take years to get to press), together with a certain apprehension (will they tear us to pieces?!), as one awaits commentary from highly respected peers. But, as the editor of this journal suggested in one of his letters, it was certainly 'worth the wait'. The commentaries on the cognizer's innards and her hypothesized capacity for representational redescription (henceforth RR) provided an extremely constructive and thought-provoking evaluation of the issues raised in our target paper. Many important suggestions were made, in particular across the various commentaries a series of very different re-interpretations of the empirical data we presented, but to answer them all would require another full-length article. In our response, therefore, we have limited our comments to a number of recurring main themes, as follows:

- (1) Is RR about representational change or process change? In modelling RR, should change occur at the level of hidden units and/or connection weights?
- (2) Can deeper understanding and further explorations of relatively simple 'first order' networks capture the intuitions underlying RR, without calling on 'second order' or hybrid models?

As was the case for the target article, both authors contributed equally to this commentary. Address for correspondence: Andy Clark, Department of Philosophy, Washington University, One Brookings Drive, Campus Box 1073, St. Louis, MO 63130-4899, USA. Annette Karmiloff-Smith, MRC-CDU, 4 Taverton Street, London WC1H 0BT, UK. Email: andy@twinearth.wustl.edu; annette@cdu.ulc.ac.uk.

- (3) Is copying/skeletonization a good candidate for simulating RR?
- (4) To what extent are nativist ideas compatible with the connectionist framework?
- (5) Are non-linguistic redescription simply an extension of the redescriptive processes enabled by the acquisition of a natural language?
- (6) What role do external notations (writing, diagrams, etc.) play in the development of a flexible mind/brain?
- (7) Is RR always failure-based and exogenously triggered, or can RR also at times be endogenously driven? What do we mean by 'external pressures', and what are the constraints on RR?

1. Representational Change versus Process Change

Several commentators raised questions about whether the account of RR is at the level of claims about representation or process. The processes that generate redescription and the different resulting representational formats are both discussed in detail elsewhere (Karmiloff-Smith, 1992a). But let us reiterate that our focus here is predominantly on the representational level and that the story remains uncommitted with respect to the issue of mechanism. We concede that our original talk in terms of 'data structures' gave an unnecessarily static and mentalistic impression. Whilst the target paper has little to offer in the way of speculations about mechanism, we are convinced that one of the attractions of the connectionist framework in contrast with classical symbolic frameworks is the fact that the processing/representation distinction is blurred.

Bechtel couches his questions about process and representation in terms of whether it is the pattern of connection weights or the pattern of activations on hidden units that might be at issue. We would reiterate, however, that for the moment RR is pitched at the level of different formats of representational change rather than the processes involved in such change. For all we know, both long-term weights and short-term activation profiles may, on occasion, be the objects of (different) redescriptive processes. These are issues for future research.

In common with several other commentators, Abrahamsen insists on the importance of attending more closely to the issue of the potential mechanisms of redescription. She offers an alternative interpretation of the data we used to support RR, suggesting that new processes may be added to generate flexible utilization of old representations rather than adding new representational formats. But, as Abrahamsen herself acknowledges, such a solution would imply a stage-like view which recent research in developmental cognitive science, particularly in infancy, seriously challenges. Abrahamsen's alternative might be more parsimonious in some domains of learning (e.g. drawing which, because of its external trace, may not have been the best arena for exploring internal representational

change—see discussion in Karmiloff-Smith, 1992a, Chapter 6). However, doing away with the process of representational redescription may be constraining with respect to other domains. Some processes can only receive specific kinds of inputs. Thus, if knowledge is stored in a particular representational format (e.g. analogue, kinaesthetic, etc.), its use may be constrained to that format and not be useable by other processes unless re-represented in a suitable format. Where there are overlapping features between domains, as in one of the examples discussed by Plunkett, there may be no need for redescription. But where the representational relations are not in the same format, then, in order to build flexible and new relationships across the representations from different domains, the representations supporting particular domains ultimately have to be re-described into a similar representational format. In any case, we should not exclude the possibility that both representational change and process change occur during development. Abrahamsen's suggestion, that processes may have their origin in procedures and become domain-general, only by a process of development that reaches a few domains in succession before generalizing to all well-developed domains, is worth pursuing further.

An issue raised in several commentaries concerns the ways in which networks might progressively change their own architectures. One recently explored solution is the use of neural networks whose connectivity and certain other parameters are controlled by genetic algorithms (Parisi, 1990; Sales, in progress). Others have explored interconnecting special-purpose modular networks, a solution to which both Abrahamsen and Bechtel suggest giving thought (e.g. Jacobs, Jordan, Nowlan & Hinton, 1991). Clark (1993) has explored this in his discussion of a network (Jacobs, Jordan & Barto, 1991) which begins life with an architecture already comprising a set of distinct modules but which does not know, at the outset, how to use those modules in a solution to a complex problem (*i.e.* it does not embody task-specific information, although of course it could if the modules were especially designed for a given task). The overall architecture then learns how to divide up the problem into subtasks, and the modules compete to perform each subtask. A beneficial result of this is that if the modules are in any way different from one another (e.g. one has more units and another less, or one computes a linear function and the other a nonlinear one), then the module best suited to that subtask will win out.

Such a process has several advantages of relevance to our overall discussion. An architecture that has assigned subtasks in this way allows useful transfer of learning since, when faced with a new but related problem, the architecture will assign similar subtasks to the same module, which will learn faster as a result. It is also protected, to a degree, from the unwelcome effects of unlearning, since radically different functions will not be assigned to a single module (see Jacobs *et al.*, 1991, p. 223). The biological advantages also include reductions in the overall number

of units and in the lengths of the connections needed—crucial factors in determining neural plausibility (*ibid.*, pp. 225–6). Indeed, the process resembles the initial overproduction, subsequent pruning and progressive parcellation/modularization of neuronal circuits in the developing infant brain (see Johnson & Karmiloff-Smith, 1992, for discussion). The removal of redundant hidden units by pruning has been shown to improve generalisation (Thodberg, 1993).

Jacobs *et al.* refer to this as a process of self-modularization. However, we would draw a distinction between what is in essence a process of self-determined exploitation of *existing* modules, on the one hand, versus progressive self-modularization, on the other. It is the latter of these two processes (whether the modules can be progressively formed from scratch and/or from very skeletal starting points—see the discussion of gradual modularization in infancy in Karmiloff-Smith, 1992a) which may shed light on the potential mechanisms of representational redescription. Suppose, for instance, that initial learning in a domain involved training a single net to perform the task. The solution would be efficient but task bound, incapable of having its elements exploited separately farther afield. Endogenous pressure could cause the brain to then use that single net as the source of a training signal to an architecture (such as Jacobs *et al.*'s) in which several modules compete and perform subtasks. The successful training of such an architecture would allow the wider and more flexible deployment of the knowledge, and the tendency to assign similar elements of future tasks to the appropriate existing module would encourage increasing informational integration over time. Such a process yields a vision of change in which the product is indeed not a classical data structure, but rather a set of mini-processors adapted to deal with the sub-elements of a task (for further discussion, see Clark, 1993).

But, in concluding this section, we agree with Bechtel that we do not yet have a fully adequate vocabulary for describing the more complex processing that may be involved in RR. Several different types of accounts have been explored, and we still feel that the connectionist framework offers the most promising path.

2. RR and Exploring Relatively Simple First-order Networks

Plunkett's basic premise is that what we termed 'first-order connectionism' may have more resources for modelling episodes of redescription than we imagined, and that as a research strategy we should push simple networks to their limits. He speculates that the process of training a single net on multiple tasks may yield interference effects which 'reach back' and change previously successful behaviour into a new form (*e.g.* his reanalysis of the French determiner example in terms of quantifiers). We concede that this is, at times, a plausible scenario, although other data from the same set of studies show how the partitive is also made use of in a non-

quantifier way (e.g. 'même' and 'même de' used at a certain period of development to mean respectively 'same-identity' or 'same-analogy'). In other words, in contrast to the RR account, Plunkett's solution may require *ad hoc* explanations for each of the developmental phenomena. Moreover, the other quantifier expressions Plunkett refers to (e.g. 'beaucoup de') are not only in the adult input, but also much earlier in the child's output than the surface over-markings for 'un de X'. Thus, if Plunkett were right, then the representational space concerned with quantification should have emerged much *earlier* in the weight matrix, causing the surface over-markings to be produced considerably earlier than the data actually demonstrate. This serves to reiterate Abrahamsen's important point that the timing of RR (as well as general constraints on RR—a point to which we return in a later section) must be a crucial component of the model.

Bechtel also argues in favour of a relatively simple network solution to the RR intuitions. He suggests that knowledge representation might be captured on a network's hidden units by adding an additional set of output units which would explicitly represent the similarity clusters that the hidden units would develop through a competitive procedure. This interesting suggestion will have to be explored within the cautionary note raised by Abrahamsen regarding the value of cluster analyses.

We found Plunkett's discussions of temporal parity and the past tense/plural net particularly useful. Connectionist simulations have certainly convinced us that full behavioural mastery may not be a prior requisite for the processes of change to start to take place, for indeed one can detect at the hidden layer representations that are not yet apparent at the output layer. For real children, whose brain functioning we cannot inspect in the same way (although technological advances in Evoked Response Potentials as well as structural and functional Magnetic Resonance Images may make that possible), we are dependent on observations at output level.

We fully endorse Plunkett's observation that explicitness, being tied to the wider exploitability of information, names a continuous property and not a simple dichotomy between implicit and explicit knowledge. But RR has always argued against a mere dichotomy. That is why we maintain, contra Scutt and O'Hara, that a developmental dimension *is* essential to understanding the human mind, since studies of adult cognition could indeed lead us to a dichotomous (implicit versus verbally explicit) view of the mind. What is new in Plunkett's formulation is that RR may turn out to be more context-sensitive than originally thought (see discussion in Clark (1993), Chapter 6). In response to both Bechtel and Plunkett, we would stress that it is the top end of this continuum of explicitness which most strongly moves us to look beyond first-order connectionism: at this top end, acquired information is exploitable not just to serve one or two different goals (as in Plunkett's past tense/plural case) but to serve a relatively open-ended set of problem-solving projects. It is this fuller flexibility (corresponding to beings who satisfy Evans' generality constraint) which, we still believe, is only made possible by courtesy of the processes of representational redescription.

3. *RR and Skeletonization*

Several commentators deemed our RR-inspired discussion of Mozer and Smolensky's technique to be interesting. However, Bechtel in particular questions how the copying and skeletonization technique change anything in the explicitness of the ensuing representations. We must agree. As he points out, amongst all the skeletonized nets, how does the system know which net to use as a template for a new problem? And how could the technique work in cases where the form of the input (e.g. number of units) differs from problem to problem? A further qualification comes from Abrahamsen. As she points out, the skeletonization procedure is limited in scope since the gross form of input/output encoding would need to be invariant across the domains (though see Plunkett's discussion of the past tense/plural case for an example in which our procedure would be workable; but that is because, as he himself acknowledges, there are many overlapping features between the two domains). One attractive suggestion made by Abrahamsen is that we might overcome some of the above problems if the copying involved a specific path through a network (perhaps tagged as to its successful function, as suggested by Bechtel), rather than a copy of the entire network. In any event, we hope that the various alternatives offered by our commentators will spur Abrahamsen's 'engineer-readers' to explore computationally the concepts that are being discussed from our 'priestly'¹ heights!

4. *Connectionism and Nativism*

Scutt and O'Hara nicely describe part of our underlying motivation as the search for an epistemological stance that stresses a third source of development, *viz.*, knowledge neither innately specified nor directly acquired by environmentally-driven learning. The key to this third position, we believe, is to focus not on 'innate representational content' but on some kind of innate 'information management', *i.e.* innately specified architectures and/or computational mechanisms (see discussion in Karmiloff-Smith, 1992b, in Ramsey & Stich, 1991, and in Elman *et al.*, forthcoming). Although connectionists have adopted the simplifying research strategy of positing initial *random* weights and connections, we see no contradiction between nativist underpinnings and connectionist-type learning. There is surely more learning in early infancy and more structure in the environment, than previous developmental theories ever

¹ We cannot resist an anecdote here. AK-S's daughter told her mother one day that there was one way in which she had failed her. Devastated, AK-S thought 'not enough love, not enough support, not enough freedom', etc. Yara retorted 'not enough religion!' So AK-S would like to thank Abrahamsen profusely for finally elevating her atheist efforts to priestly endeavours!

suspected. The connectionist framework is in our view an ideal lens through which to explore the interplay of different types of developmental predispositions of a representational, architectural and/or computational type, depending on the domain.

5. The Role of Language

Dennett seeks to convince us that the driving force behind the decoupling of acquired knowledge from its task-specific roots is not primarily an internal, non-linguistic process of redescription, but is instead the effect of the child's internalization of the structures of public language and culture into its own brain. His claim is that redescription is a capacity which flows from the antecedent (or, at least, simultaneous) mastery of language.

We gave some reason, in our original text, for not endorsing such a view. Dennett maintains that those particular reasons were not conclusive, and we agree. So in defence of our position, we will add some more recent considerations drawn from studies of prelinguistic infants and of certain abnormal phenotypes (Karmiloff-Smith, 1992b). But prior to that, let us reiterate the view we still hold of the role of language in redescriptive processes.

Language users, we agree, are the most profound and far-reaching redescribers of their acquired knowledge. We take the linguistic format to constitute one of the uppermost levels of redescription (what Karmiloff-Smith, 1986, 1992a, called 'level E3' which, for space reasons, we collapsed with level E2 in the target article and called 'levels E+'). Linguistic representations play an important role both within and between individuals. First, within the cognizer, linguistic redescription is the one representational format that frees knowledge from spatial, temporal and causal constraints and enables new links to be noted across originally different representational formats (e.g. spatial and kinaesthetic). Second, redescribers who can speak a public language can also exploit the redescriptive successes of other agents. In other words, the benefits of RR within an organism are enhanced by interactions between organisms (and across generations) using the same public language. We agree with Dennett and Bechtel that acquisition of a public language extends the nature and range of a being's redescriptive skills in important ways. But as to whether language is itself the root of the basic mechanisms of redescription, we remain unconvinced, for the following reasons.

First, *prelinguistic* infants have been shown to go beyond the knowledge embedded in their special-purpose procedures for responding to environmental stimuli (Mandler, 1988, 1992). On the basis of a number of infancy studies, Mandler used the RR thesis to argue that infants too redescribe their early procedural representations, not into propositional form, but into analogue form. Such analogue representations are still relatively primi-

tive but, according to Mandler, they are declaratively accessible to the infant, and embody concepts such as target, source, path, trajectory, direction, etc. The analogue representations form the transition between procedural representations and the subsequent linguistic ones. Thus, we question both Dennett's contention that 'concepts are things in our world because we have language' and Bechtel's that 'it is through learning to comprehend and produce natural language utterances that cognition acquires the degree of productivity and systematicity that it exhibits'. Although language clearly enhances cognition, much developmental research has shown that prelinguistic infants develop a rich variety of concepts explicitly defined in their internal representations (Mandler, 1992). By contrast, some linguistically-sophisticated but abnormally developing children and adults are surprisingly lacking in conceptual development, a point to which we now turn.

Indeed, our second reason for maintaining our original view about the role of language stems from recent work with two abnormal phenotypes (Karmiloff-Smith, 1992b). Williams syndrome, as well as hydrocephalus with associated myelomeningocele, can result in an uneven cognitive profile in which language is one of a small number of spared domains in the face of other serious deficits. There are children, adolescents and adults who have IQs in the 50-70 range, who have extreme difficulties with arithmetic, reading, spatial tasks and problem solving, and yet who have large vocabularies, good syntax, and often speak very fluently and engagingly.² Most striking is the fact that children with these two disorders often fail some simple problem-solving tasks that chimpanzees can solve (Volterra, personal communication). Moreover, despite their fluent language, even in adolescence and adulthood these subjects build much weaker implicit systems and fail to build the explicit explanatory theories typical of normally developing children (Carey, 1993; Karmiloff-Smith, Grant & Berthoud, 1993). Yet they are extremely good at learning new and unusual linguistic labels; in fact they are significantly better than normal 5-year-olds (Karmiloff-Smith, Grant & Berthoud, 1993). Thus the acquisition of a natural language can occur without the operation of the processes of representational redescription that make possible the flexibility and creativity of normal human cognition. The language of these abnormal populations seems to be a special-purpose syntactico-lexical device. In our view, something other than linguistic labelling underpins RR. In other words, a crucial distinction must be drawn between linguistic representations in their communicative function, in which linguistic labels may be relatively shallow in their cognitive underpinnings, and linguistic *re-representations* which lead to ever-increasing relations between different aspects of the cognitive system.³

² We are not talking here about cases of so-called empty 'cocktail speech'.

³ For space reasons, here we deliberately leave aside consideration of the further issues raised by the so-called Baldwin effect (see Dennett's closing comments).

6. *The Role of External Notations*

First, two cautionary notes. In considering the role of external notations, we must be careful to avoid slipping between the on-line, fast-fading message of spoken language, on the one hand, and external notations such as written language which leave a trace and can be contemplated off-line, on the other. Second, unlike Dennett, we argue that the external traces left by other species are non-intentional in any interesting sense of the word (see discussion in Karmiloff-Smith, 1992a, Chapter 6). We fully agree with Dennett and Bechtel that we need to explore the various ways the cognitive load may be divided between innards and the external environment (for studies of children's invention, use and understanding of external notations such as diagrams, maps and written notations, see Karmiloff-Smith, 1990; Lee & Karmiloff-Smith, 1992; Tolchinsky-Landsmann & Karmiloff-Smith, 1992). Bechtel argues that where we see an internal process of redescription, the real story may be better described in terms of 'changing interactions between a cognitive system and external representations'. It is indeed true that learning to read, for instance, makes possible new representations of the structure of words in the lexicon, in that phonemic (as opposed to syllabic) components can now be extracted. Illiterate adults do not have a phonemic level of representation (Morais, Cary, Alegria & Bertelson, 1979).

We might again raise a question similar to that discussed in the previous section: Is the use and understanding of external notations dependent on language? Take Dennett's interesting discussion of the use of arbitrary external markings as a problem-solving aid in a hiding game. First, it is to be noted that prelinguistic infants can successfully use such marks in a hiding task (Piéaut-le Bonniec and de Schonen, 1976a, 1976b). Second, even fluent-speaking Williams syndrome adolescents capable of counting to, say, twenty have extreme difficulty in making use of a coloured marker placed on one of seven counters arranged in a circle, to know where to begin and end their counting sequence (Karmiloff-Smith, 1992b), whereas normal 5-6-year-olds find this easy (Gold, 1987). Thus, while we agree that external notations are an important enhancer of representational flexibility and manipulability, we stress that if the use and understanding of external notations changes a cognizer's internal representations, then something must surely drive and enable that change. That something could be the availability of new inner resources whose target objects are external representations. Such a vision would represent a shift to a more process-oriented conception of the redescriptive processes (see also Clark, 1993).

7. *Exogenous versus Endogenous Change, and Constraints on RR*

The description of RR as 'spontaneous' caused some difficulties for a number of commentators, particularly Plunkett, and Scutt and O'Hara. Of

course, as Scutt and O'Hara suggest, on any coherent reading of RR, there must be pressure for redescription from *somewhere*. Hence it cannot really be spontaneous. To set the record straight:

- (1) We should situate the beginnings of the RR hypothesis (Karmiloff-Smith, 1979) in its historical context. At that time, only external, failure-based models of learning were given consideration. All learning studies were based on negative feedback, *i.e.* the change in the learner's behaviour on the basis of failure to reach a goal. So K-S felt obliged to lay great stress on the fact that RR *could* also be internally and spontaneously generated. Now other developmentalists endorse the possibility of success-based learning (Siegler & Munakata, 1993).
- (2) It has never been claimed that redescription is causeless: only that it did not require prompting by gross failure to succeed at some external task.
- (3) Even this was not intended as a *necessary* condition of a process being properly described as redescriptive. Of course gross failure *can* prompt redescription, but we argue that redescription can also occur without any such external prompting.

In common with RR, Plunkett stresses an account that calls for a reorganization of existing knowledge not prompted by gross failure on tasks involving the knowledge in question. But more in the spirit of Piagetian theory, Plunkett locates the source of the reorganization in the system's internal attempts to accommodate new learning. However, where Piaget invoked internal *conflict* across different systems, we, like Plunkett, see the processes of change more as the result of the internal system dynamics (see Karmiloff-Smith, 1992a, for discussion). None the less, it is worth noting, in this context, that we accept Plunkett's point that, if we are thinking in connectionist terms, it could easily be that a small internal error signal persists *despite* gross, external success and that the attempt to reduce this residual error could constitute an internal 'failure-driven' spur to redescribe, *i.e.* to reorganize a body of stored knowledge (see Plunkett's discussion of temporal parity; a case also addressed in Clark, 1993). It is also possible that external and internal pressures play different roles at different points in the redescriptive cycle (see Karmiloff-Smith, 1992a, for discussion), which once again underlines the need for much greater specification of the processing model implicit in RR, as called for by most of the commentators.

Abrahamsen and Scutt and O'Hara raise serious problems about the constraints on RR, fearing either a 'copy-happy' process or 'random-RR' resulting in no apparent gains to the cognitive system. We fully agree that the issue of constraints on RR remains an unsolved issue. We have no problem with invoking internal pressures, but the available empirical data do not convince us that RR is always generated by, say, a lack of memory

space, as Scutt and O'Hara suggest. Some explorations of the constraints on RR have been pursued, but so far with limited success. For example, it has been shown that competition between two simultaneously developing processes in the acquisition of discourse rules results in one of the processes becoming available to metalinguistic reflection and the other not, even in adults (Karmiloff-Smith, Johnson, Grant, Jones, Karmiloff, Bartrip & Cuckle, 1993). Of course, expert linguists do manage to reflect metalinguistically on both processes but, interestingly, they cannot carry out discourse analyses directly on on-line spoken language, but only when the spoken message is frozen into a static written form, *i.e.* an analysis of an external notation. This again highlights the fact that external notations extend the powers of the mind. Another attempt to understand the constraints on RR involved a pilot study of the possible role of puberty in constraining further redescrptions. But the results were inconclusive and the study gave rise to no general account of the constraints on RR.

The commentators have indeed put their finger on a particularly important problem with RR as currently conceived. Why do we not redscribe and link *all* the knowledge we store? We know that adults and children can hold contradictory theories about the world. Why is it that we often have to await university for an expert to point out to us an unsuspected relationship between two very different domains that we already know about? Why was this not part of our own spontaneous redescriptive processes? To reiterate our comments in a previous section, the role of others having created external notations of different domains of internalized knowledge may be crucial here in discovering potential cross-representational relationships.

8. Concluding Remarks

We still believe that human cognition is special, that a developmental perspective is essential to understanding its final form, and that what we would need to add to Dennett's cuckoo to enable it to exploit the knowledge embedded in its special-purpose procedures is not language *per se*, but something akin to representational redescription. Extensive training with arbitrary representational codes does seem to give rise to some rudimentary redescrptions in the chimpanzee (see discussion in Karmiloff-Smith, 1983, 1988), but RR is nothing like as pervasive in that species as it is in human development. However, we entirely agree that we are still far from specifying the precise mechanisms that would have to be built into the brains of other species. Through the various commentaries, it is clear that RR can be defined by a cluster of somewhat different ideas which must individually face the tribunals of further investigation.

We should like to conclude by endorsing a comment in Dennett's opening paragraph, in that we genuinely hope to have begun to 'transform an acrimonious and ideology-ridden debate' about rules versus no rules,

symbolic processing versus connectionism, 'into a multi-perspective research program'. Hopefully our target article, the illuminating commentaries and numerous re-interpretations of the data, together with our defence of RR in our response, will generate new philosophical, developmental and computational research.

Cognitive Development Unit
Medical Research Council
4 Taviton Street
London WC1H 0BT
and
Department of Psychology
University College London
UK

Department of Philosophy
Washington University
St. Louis MO 63130-4899
USA

References

- Carey, S. 1993: The Development of Explanatory Theories in Subjects with Williams Syndrome. Poster Presentation, Society for Research in Child Development Biennial Meeting, New Orleans.
- Clark, A. 1993: *Associative Engines: Connectionism, Concepts and Representational Change*. Cambridge, MA.: MIT Press.
- Elman, J.L., Bates, E., Johnson, M., Karmiloff-Smith, A., Parisi, D. and Plunkett, K. (in press): *Rethinking Innateness: Connectionism in a Developmental Framework*. Cambridge, MA.: MIT Press.
- Gold, R. 1987: *The Description of Cognitive Development: Three Piagetian Themes*. Oxford: Oxford University Press.
- Jacobs, R.A., Jordan, M.I. and Barto, A. 1991: Task Decomposition through Competition in a Modular Connectionist Architecture: The What and Where Visual Tasks. *Cognitive Science*, 15, 219-50.
- Jacobs, R.A., Jordan, M.I., Nowlan, S.J. and Hinton, G.E. 1991: Adaptive Mixtures of Local Experts. *Neural Computation*, 3, 79-87.
- Johnson, M.H. and Karmiloff-Smith, A. 1992: Can Neural Selectionism be Applied to Cognitive Development and its Disorders? *New Ideas in Psychology*, 10, 35-46.
- Karmiloff-Smith, A. 1979: *A Functional Approach to Child Language*. Cambridge: Cambridge University Press.
- Karmiloff-Smith, A. 1983: A New Abstract Code or the New Possibility of Multiple Codes? Peer commentary on D. Premack, Mental Codes in Man and Beasts. *The Behavioral and Brain Sciences*, 6, 149-50.
- Karmiloff-Smith, A. 1986: From Metaprocess to Conscious Access: Evidence from Children's Metalinguistic and Repair Data. *Cognition*, 23, 95-147.

- Karmiloff-Smith, A. 1988: Human versus Non-human Abilities: Is there a Difference which Really Counts? Peer commentary on Davis and Perusse, Numerical Competence in Animals. *Behavioral and Brain Sciences*, 11, 589-90.
- Karmiloff-Smith, A. 1990: The Human Print-out Facility: Extending Biological Constraints by Cultural Tools. Paper presented at the Conference on Domain Specificity and Cultural Knowledge, University of Michigan.
- Karmiloff-Smith, A. 1992a: *Beyond Modularity: A Developmental Perspective on Cognitive Science*. Cambridge, MA.: MIT Press.
- Karmiloff-Smith, 1992b: Abnormal Phenotypes and the Challenges they Pose to Connectionist Models of Development. *Technical Reports in Parallel Distributed Processing and Cognitive Neuroscience*. Technical Report PDP.CNS.92.7, Carnegie-Mellon University.
- Karmiloff-Smith, A., Grant, J. and Berthoud, I. 1993: Within-domain Dissociations in Williams Syndrome: A Window on the Normal Mind. Poster Presentation, Society for Research in Child Development Biennial Meeting, New Orleans.
- Karmiloff-Smith, A., Johnson, H., Grant, J., Jones, M.-C., Karmiloff, K., Bartrip, J. and Cuckle, P. 1993: From Sentential to Discourse Functions: Detection and Explanation of Speech Repairs by Children and Adults. *Discourse Processes*, 16, 565-89.
- Lee, K. and Karmiloff-Smith, A. 1992: Cognitive Constraints on Notations: Encoding States versus Transformations (submitted).
- Mandler, J. 1988: How to Build a Baby: On the Development of an Accessible Representational System. *Cognitive Development*, 3, 113-36.
- Mandler, J. 1992: How to Build a Baby II: Conceptual Primitives. *Psychological Review*, 99, 587-604.
- Morais, J., Cary, L., Alegria, J. and Bertelson, P. 1979: Does Awareness of Speech as a Sequence of Phones Arise Spontaneously? *Cognition*, 7, 323-31.
- Parisi, D. 1990: Connectionism and Piaget's Sensori-motor Intelligence. Paper given at the conference on *Evolution and Cognition: The Heritage of Jean Piaget's Genetic Epistemology*. Bergamo, Italy, October 5-8, 1990.
- Piéraut-le Bonniec, G. and de Schonen, S. 1976a: Etude Génétique d'Activités Sémiotiques: Utilisation de Traces et de Marques chez l'Enfant de 2 à 5 ans. I. Utilisation du Marquage d'un Object. *L'Année Psychologique*, 76, 1, 55-77.
- Piéraut-le Bonniec, G. and de Schonen, S. 1976b: Etude Génétique d'Activités Sémiotiques, II. Utilisation de Traces Laissées par une Activité. *L'Année Psychologique*, 76, 2, 401-416.
- Ramsey, W. and Stich, S.P. 1991: Connectionism and Levels of Nativism. In W. Ramsey, S.P. Stich and D.E. Rumelhart (eds.), *Philosophy and Connectionist Theory*. Hillsdale, NJ.: Erlbaum, 287-310.
- Siegler, R.S. and Munakata, Y. 1993: Beyond the Immaculate Transition: Advances in the Understanding of Change. *Society for Research in Child Development Newsletter*. Winter.
- Thodberg, H.H. 1993: Ace of Bayes: Application of Neural Networks with Pruning. *Neuroprose*.
- Tolchinsky-Landsmann, L. and Karmiloff-Smith, A. 1992: Children's Understanding of Notations as Domains of Knowledge versus Referential-Communicative Tools. *Cognitive Development*, 7, 287-300.