

Computation, Dynamics and Sensory-Motor Development*

Julie C. Rutkowska
Cognitive & Computing Sciences,
University of Sussex, Brighton, BN1 9QH, UK.

email: julier@cogs.susx.ac.uk

phone: +44 1273 678946

fax: +44 1273 671320

Abstract

Outside Piagetian theory, sensory-motor co-ordinations are often relegated to the domain of 'mere' motor skill, but their development shares important features with that of cognitive structures. This paper focusses on early action, to assess the prospects for an account of development that is based on a system's initial mechanisms and processes of interaction with its environment without prespecification of stable patterns of organization that will be acquired. A truly epigenetic account proves elusive, with empirical findings increasingly being taken to indicate preadaptation and strong domain-specific constraints on infant abilities. Despite this, evidence can be marshalled for variability, which is compatible with a general-purpose process of internally motivated and structured organizational change. However, the mechanisms underlying this process are obscure.

Clarification is sought by considering current attempts to understand sensory-motor co-ordination through the construction of artificial agent-environment systems. Disappointingly, such approaches often share a need to incorporate an explicit bias towards the recurrent behaviour patterns that will come to have functional significance for the systems they aim to explain. Synchronic systems in this vein exploit predesigned sensory-motor connections with the environment. Their diachronic counterparts feature designer specification of acceptable outcomes for activity in the form of problem-specific fitness functions or goal-like value schemes that are credited to evolution. Neither computational nor dynamical systems concepts provide an automatic escape from this problem, but most promising may be robotics approaches informed by dynamical systems theory that challenge mainstream views of information and information-processing.

*Paper presented at the symposium on *Modelling Epigenetic Emergence*, convened by Jan Boom (Utrecht) and Peter Molenaar (Amsterdam & Penn State) for the Piaget Centennial Conference, *The Growing Mind: Multidisciplinary Approaches*. Geneva, 14-18 September 1996.

1 Reviving epigenetic issues

The structures of knowledge do indeed achieve necessity: but at the end of their development without having it from the start, and do not involve any antecedent programming. (Piaget, 1972, p.56)

Piaget speaks here of formal properties of operational thought and genetic programs, but his general concerns are equally relevant to the kinds of sensory-motor co-ordinations that many have relegated to the domain of ‘mere’ motor skill, and to the role of computational approaches in explanations of development.

Increasingly, the preoccupation of much mainstream infancy research with ‘between the ears’ cognition is being challenged by the view that mind is grounded in action (e.g. Rutkowska 1993; Thelen & Smith, 1994). Whether action is seen as a precursor of cognition, or action–cognition as a mistaken opposition, action and cognition pose identical problems. Those problems mark an interdisciplinary revival of Piaget’s traditional concerns: What kind of processes give rise to developmental outcomes that are not predetermined, however predictable their acquisition appears to be? Can an action-based, epigenetic approach to development surpass and supplant inadequate nativist or empiricist accounts of our knowledge of the world?

In this paper, I shall be looking at these problems from the standpoint of (apparently simple) sensory-motor acquisitions. Two points are especially pertinent to establishing the broader relevance of this perspective:

- Acquisition of everyday sensory-motor activities meets criteria that have been proposed for strongly constrained knowledge structures, and taken to support Chomsky’s nativist view of natural development as a form of growth that is guided to a predetermined end by domain-specific preadaptations (Keil, 1981). Activities such as locomotion and prehension exhibit mapping from a wide range of experience onto a narrow range of outcome structures; they appear to be rapidly, universally and effortlessly acquired without formal tutoring; the psychological mechanisms underlying them are relatively closed to introspection; and there is a strong sense of anomaly with violations (Rutkowska, 1991, 1993).
- Despite this apparent universality, evidence for variability in the form of developmental acquisitions questions the extent to which they are prespecified. Flexibility is exhibited at the level of movements, as when locomotion is attained by sitting cross-legged and using the arms to pull the body along in ‘scotting’ (Dennis, 1960); body parts, as when prehension via legs and feet routinely replaces arm–hand use in differently-abled thalidomide subjects; and sensors involved, as when feeling a surface substitutes for seeing depth in control of crossing behaviour on the visual cliff (Rader, Bausano & Richards, 1986).

The explanatory framework that I shall explore features recent computational work that shares developmental psychology’s growing focus on action and emphasises the necessity for models of whole agent–environment systems. The attempt to understand developmental processes through computational ideas is far from new. In 1962, Simon

suggested writing a general-purpose learning program that would operate on other programs to generate the kinds of organizational change that Piaget attributed to equilibration. But although Simon's notion may be a useful metaphor for describing a system's potential for adaptive change, the extent to which it can explain such change is hampered by its use of classical computational concepts in a traditional way, especially the idea that change must implicate a distinct mechanism component that is dedicated to 'doing the changing'. This is precisely the kind of idea that recent computational work challenges. In the next section, I briefly outline some basic assumptions of this work by locating them in their historical context. I also draw attention to what I consider to be a persistent source of difficulty for our attempts to frame novel theories from computational or psychological directions.

2 Self-organization and the 'R-words'

Computational approaches associated with traditional cognitivism have predominantly attributed intelligent systems' knowledge of the world to internal representations of an objective external reality, viewed as symbol structures that make explicit information about objects, their properties and their location with regard to one another and to the subject. Programs, which specify rules for manipulating these structures, govern the reasoning processes for formulating goals and planning behaviours that are identified with the core of the mind's activity. So, better (explicit, exhaustive) representations should yield more knowledge; and better (general) programs should yield more ways of deploying that knowledge. In practice, systems built along these lines prove to be notoriously brittle — a system may be good at a game like chess, but will be stopped in its tracks by encountering even another game environment. It seems virtually impossible to get into a single system all the knowledge and program rules that seem necessary for flexible, adaptive behaviour.

Subsequent computational work in parallel distributed processing or connectionism aims to improve on the traditional computational picture by characterizing cognition at a subsymbolic level, in terms of multiple, interconnected simple units operating in parallel (cf. neurons, though how appropriately is debatable). Computation as fixed, sequential programmed rules is replaced by the whole network settling into a stable pattern of activation by trying simultaneously to satisfy many soft or weak constraints that are only meaningful if considered collectively. Compared with traditional computational systems, connectionist networks have achieved relative flexibility on some tasks under noisy or variable circumstances by settling into the most likely of a range of related solutions, with pattern discrimination or categorization considered one of their great implementational successes. Two (attempted) shifts are in evidence here:

- Increasing concern with emergent properties. For example, rule-like behaviour is attributed to patterns of activity that recur within a network, not to operation of a pre-existing program.
- Increasing exploitation of self-organizing properties. Notions like 'settling into a solution' move the focus from program design towards interest in how systems gen-

erate and maintain their own organization (i.e. are ‘self-producing’ or ‘autopoietic’ in Maturana and Varela’s (1988) sense).

Computational work that focusses on whole agent–environment systems aims to go even further in these directions. Its emphasis is on putting action and cognition into context, as processes of physically embodied systems that are embedded or situated in an environment. Classical systems mistakenly used the ‘in principle’ separation of program and physical machine to licence total disregard for the physical circumstances of the cognition they attempted to model. Connectionist systems too are limited. Despite pleas to biological plausibility, they remain far from modelling real-life deployment of mental processes or development. Their sensory interfaces with environmental inputs rarely consist of intensity arrays, tending to involve experimenter selection and hand-coding to a degree that questions the label ‘self-organizing’; networks generally model or simulate only isolated subsystems, rather than being part of a whole system that is embedded in a real environment; and their changing organization relies extensively on explicit external feedback about the correctness or otherwise of outputs. In a significant sense, their designer’s/experimenter’s activities constitute their environment, closing the loop between their input and output units. By way of contrast, the autonomous agents direction aims to shift the research focus more strongly towards self-organization, by building robots that are left to their own devices in real environments (or at least simulating such complete agent–environment systems), and by capitalizing on the unique forms of emergence that this embedding of cognitive and physical processes may support.

A key success of this new computational direction has been its contribution towards dismantling the Cartesian mind–body dualism that pervades much of cognitive science. But there are many dualisms. Costall (1995) aptly appropriates one of Benchley’s aphorisms — that there are two kinds of people in the world, those who believe that there are two kinds of people and those who don’t — for the psychological case: “the world can be divided into two kinds of psychologists — those who know they are committed to some form of dualism and those who don’t.” The insidious dualism to which I wish to draw attention here is that of subject–environment.

We should make this distinction far less readily than we do. As observers, we generally make it in physical terms; skin or tin boundaries demarcate the limits of the subject in their environment. What is needed is a good process-oriented language that spans the joint roles of subject and environment in generating functional activity. What we predominantly have is a resolutely linear, directional, in–out view of environment’s relation to mind. Senses and motors are physically/spatially separable, interfacing with the environment in different places, hence the processes they implement are separated too. It is remarkably difficult to escape the worldview that sees the senses as the front-end of any intelligent system, with its motor capacities bringing up the rear as its output. A system uses its senses to get information from the environment, then determines how it will act. Constraints on adaptive action are (mis)localized, in either internal cognitive structures or external task–domain ones.

This directional, in–out treatment of sensory-motor and cognitive functioning supports not only the traditional nativist–empiricist split but also contemporary debates about the relevance of representation to valid accounts of adaptation and intelligence. I have

coined the expression ‘the R-word’ to capture the emotive tone of a good deal of the anti-representational discussion coming from new computational approaches. (And discuss elsewhere how these objections are appropriate for re-presentational mechanisms that substitute for the environment but irrelevant to action-based mechanisms that support representation by selective correspondence; e.g. Rutkowska, 1993, 1994a & b.) What seems to be missed, however, in eagerness to emphasise the significance of environmental embedding for effective functioning, is the implication of implicitly endorsing another ‘R-word’: realism, which is behind assumptions that environmental information can replace representation. For example, connectionism’s new look for cognition is not so radical as to move away from divisions into input, output and intervening units, or to move beyond ‘recovery’ or ‘discovery’ metaphors for the subject’s relationship to information. Even highly influential whole-agent research such as Brooks’s (1991) ‘intelligence without reasoning/representation’ robotics’ approach assumes that animals sensors “extract just the right information about the here and now around them.”

If our aim is a genuinely epigenetic framework, then pleas to a privileged precursor of the subject’s knowledge in external environmental information are no improvement over allocating this privileged status to the subject’s internal (model-like) representations. A key dimension of epigenetic explanations, as viewed from the vantage point of Varela’s (1988) enaction framework, is that the subject’s world is ‘brought forth’ through a history of structural coupling between organism and environment, not pre-given in one or other component of this system. For example, Varela contends that information is the phlogiston of cognitive science, repeatedly invoked as a source of pre-given order outside of the subject’s activities.

Getting to grips with emergent phenomena in action entails moving beyond our entrenched ways of considering the subject–environment relationship. Action needs to be treated as a systematic concept that refers to functional co-ordination of sensory and motor processes in the environment, not to one bit of the operation of this subject–environment system (e.g. isolated motor processes or overt behaviour; Rutkowska, 1993). The following sections of this paper follow up this line of reasoning by looking at key aspects of agent–environment systems that start out with unbiased sensory-motor connections. Such systems are often based on the assumption that human design of effective sensory and motor connections is too hard to succeed at any but a trivial scale, and that developmental/evolutionary techniques must be tried instead (see Rutkowska (1995) for comparison of these strategies). Two issues are addressed: Can functional sensory-motor connectivity be achieved by such systems? What have they acquired once they’ve achieved it?

3 Evaluating ‘value’

An important example of a self-organizing agent–environment system that aims to clarify developmental concerns is the Darwin III robot (Edelman, 1992; Reeke, Finkel, Sporns & Edelman, 1990). Its underlying commitment is to establishing the power of self-organization as a developmental framework, and to challenging Cartesian dualism. While most implementations of Darwin III feature simulations rather than a ‘real’ robot,

they nevertheless incorporate consideration of neural, behavioural and environmental contributions. Categorization is seen as the focus of getting to know the world. Categorization, however, is not seen traditionally as a central process of detecting properties of an objective external environment, as it is with classical computational, information-processing approaches. Rather, it is a sensory-motor construction — “a behavioural act.” The theory of neuronal group selection proposes that a successful adaptive system will start out with unbiased connections between groups of sensory and motor neurons. Individual history grounded in variation in experience can then give rise to ontogenetically determined ‘wiring’, from which emerge action-based categorization, knowledge and consciousness.

Say we look at the progress of a simulation of the oculomotor system, reach system (a jointed arm), tactile system and an embedding environment that contains (initially only from an observer’s perspective) an object that could be manipulated. Before training, Darwin III’s arm movements are appropriately described as random flailing; there are no recurrent sequences of movement that we, as observers, would identify with behaviour patterns. After training, however, such sequences are clearly evident. While precise movements vary somewhat from trial to trial, they clearly converge on what looks like a behaviour pattern of reaching to the object. A category along the lines of ‘graspable thing’ has been constructed through a process whose outline form bears similarities to Piagetian notions of functional and recognitory assimilation.

Darwin III’s key achievement can be characterized as discrimination between its adaptive and chance activity. Some patterns of ‘neuronal’ connectivity get strengthened (i.e. are ‘selected’ for their adaptive potential, by analogy with Darwinian evolutionary theory); others do not. How does the training achieve this? Crucial to its success is the fact that certain sensory-motor configurations are the focus of inbuilt value schemes that are assumed to be a product of evolutionary experience. In the case of reaching, for example, the system incorporates a value scheme that treats the hand being in proximity with the object as ‘good’; when this state is attained, the sensory-motor configurations that arrived at it are strengthened in a version of positive reinforcement. Sensory-motor configurations that do not result in this outcome are not negatively reinforced, but they are ignored. If the value scheme is disconnected, the system does not learn more slowly or more uncertainly, it simply does not acquire the behaviour pattern of reaching at all.

One objection to this agent–environment approach might be that it suffers from connectionism’s restricted-input problem, insofar as agent and environment are so selectively modelled that it would be surprising if reaching to the object did not get established. In fact, even in this impoverished version of a whole world, no organized behaviour gets set up at all without the value scheme. The role of the value scheme is akin to the experimenter-determined feedback on the outputs of connectionist learning networks, such as back propagation, which is used to determine which input–output connections are to be strengthened. It also resembles the role of designer-specified fitness functions of genetic algorithms research, which are generally employed for evaluation, hence selection, of sensory-motor networks that evolve in the course of trying to attain some task-specific function. Interestingly, the use of evolutionarily inspired genetic algorithms to evolve artificial neural networks for sensory-motor co-ordination may be accompanied by denials that real evolution is an orthodox process of optimisation, or that contemporary

animals can be seen as solutions to problems posed in their species' distant evolutionary past (Cliff, Harvey & Husbands, 1994). Nevertheless, reservations about locating the form of individual performance too much 'in the genes' are accompanied by continued use of fitness functions for the pragmatic purpose of getting the acquisition process to work.

Do value schemes of this kind constitute a vestigial 'ghost in the machine'? Their dominant role raises a number of issues:

- **Inbuilt goals?** Such value schemes share properties of the traditional goals of centralized, classical artificial intelligence. An observer's description of the task that the system solves is incorporated as a functional component of the agent's mechanisms. Unlike traditional goals, the value scheme does not play a role in selecting and planning the activities that will lead to the outcome it specifies. Like traditional goals, however, it provides a 'stop rule' that specifies when activity has achieved a more or less stable end-state that is deemed advantageous for the system. This looks a lot like predetermination of developmental outcomes.
- **Buck passing to evolution?** This framework places its emphasis on individual history, but evolution, in the guise of the preadapted value scheme, may be doing more of the work than an epigenetic perspective onto the system might hope for. This leads back to the problem of how a value scheme might itself evolve.
- **Acquisition or tuning?** Darwin III's designers accept that the system is not exhibiting development in the sense of radical qualitative restructuring, or possibly too much by way of learning, insofar as new values for behaviour cannot be acquired. They stress, however, that it is being trained. But how different is this acquisition of sensory-motor co-ordination from a system that starts out with (evolutionarily) biased sensory-motor connectivity then requires individual experience primarily for improvement on initially imperfect movement execution?

Records of Darwin III's training are not unlike those of changing pecking patterns in chicks (Hess, 1956). Initially, chicks' pecking is widely distributed, rarely making contact with grain; over the course of a couple of weeks, their ability to achieve sensory-motor concordance improves, so that their pecks converge to a point on/about the 'target' grain. If, however, chicks are equipped with displacing optical prisms shortly after birth, their pecking still shows increasing convergence, but before and after this refinement the locus of pecking is laterally displaced in relation to the actual position of grain. In this case, individual experience may be capitalizing on sensory-motor co-ordinations rather than establishing them *ab initio*, an interpretation that fits with the fact that pecking exhibits discrimination of the specific shape of grain-sized objects from the outset.

- **Restrictive semantics?** Potential flexibility in outcome behaviours would seem to be an important advantage of initially unbiased as opposed to biased sensory-motor connectivity. Chances of novel or unexpected acquisitions are, however, overly restricted by value schemes that have a clear semantics in terms of behaviour patterns, and which go as far as specifying the body parts involved in cases like the

prehension value scheme outlined above. While some form of reinforcement-like value may be essential for the developmental process, can it be as behaviourally transparent as this example? There surely cannot be value schemes for every recurrent behaviour pattern that infants come to display? This point is illustrated further by looking at how value schemes support a real robot's learning to discriminate between objects in its environment that are graspable, pushable or best ignored in terms of its physical-motor structure (Scheier & Pfeiffer, 1995). Not only is it necessary for objects of different sizes and weights to permit different activities on the part of the robot, grasping and pushing must be associated with explicit value schemes that reinforce the sensory-motor configurations that produce them. There is no chance that behaviours like slapping or shaking, which readily come to be part of human infants' object manipulations, could find their way into the robot's repertoire, even though it is physically capable of executing them and they might occur through random movement variation.

Overall, the status of internal value schemes raises doubts about how far the acquisitions of self-organizing systems that exploit them are emerging from a truly epigenetic process. However, the workings of such systems may be more unambiguously successful in clarifying an epigenetic perspective on what they have acquired once they become established.

4 From invariants to covariation

A traditional interpretation of what is being acquired by systems that succeed in getting their sensory and motor processes into functional alignment would consider the object recognition that is suggested by their appropriate activity as an aspect of a distinct sensory process, in keeping with the kinds of dualist subject–environment distinctions that were noted earlier in this paper. Darwin III, for example, might be considered to learn what (input) sensory patterns provide environmental information about things for which reaching is an appropriate (output) motor process to initiate. However, robotics work within the whole agent–environment system framework is beginning to offer us a far more mutual, co-relative view of how sensory-motor systems work and mesh with the environment.

Of greatest interest is robotics work that forefronts a dynamical systems analysis of sensory-motor systems. This makes it possible to see how information is an emergent property of the dynamics of subject–environment interaction — not in the trivial sense that activity is needed to detect/recover/select information, but in a more strongly constructivist sense of epigenetic emergence through action. Here, I shall not be considering whether autonomous agents are better viewed as dynamical systems than as computational systems, nor whether computational systems are a particular kind of dynamical system. In important respects, theory and metatheory are separable; for example, you do not need to commit yourself to the dynamical systems approach to assume the centrality of process issues and of a systematic perspective onto the subject–environment relation (cf. Beer, 1995). However, the dynamical systems focus on central nervous system, body and environment as variously coupled dynamical systems may have the edge

in getting at the role of temporality and ongoing history in situated systems through a process language that promises finer-grained temporal analysis than the more molar procedural notions of computational analysis.

One example features evolution of a sensory-motor controller in a recurrent dynamical artificial neural network, enabling a robot to find its way to the centre of a circular arena and to remain there (Husbands, Harvey & Cliff, 1995). There turns out to be no useful characterization of how the robot performs in terms of its sensors coming to detect an invariant property of stimulation associated with task solution, e.g. the ratio of wall height to floor radius specified by the absolute value of inputs to the ‘eyes’ at the centre of the arena.

Reverse engineering to clarify what connections have been established reveals nothing like the neat distinction between input, output and intervening units that typifies connectionist networks. No sensory and motor subsystems are found. Internal structure looks more like a spaghetti junction, suggesting that the sensory and the motor exert reciprocal influences on each other at all stages of functioning. There is no psychologically meaningful decomposition in terms of traditional information-processing components. Nor is there evidence for any component(s) that might function as a ‘smart machine’ (Runeson, 1977), more in keeping with the theory of direct visual perception, operating as a special-purpose system dedicated solely to detection of a particular invariant in the ambient optical array that can control behaviour. To the extent that such invariant detection might be considered to occur, it is implemented in the activity of the entire robot.

The implications of such findings are clarified through a dynamical systems analysis of wall-following robots. Smithers (e.g. 1994, 1995) work illustrates how improvement in performance is disappointing if you try to achieve it by giving the robot a ‘better’ set of mappings between its sensor readings and its motor acts. It is far more effective to change the dynamics of the robot’s sensory-motor functioning. For example, the time the robot spends on activities such as turning or reversing can be made proportional to the amount of sensor signals that have been received over a particular period of time. In a very important sense, recent history is made to count.

A most interesting feature of this type of organization is that the robot reveals an important property of self-organizing as opposed to rule-following systems: when it is in what looks to an observer like the ‘same’ situation, its sensory-motor configuration is rarely identical, despite the fact that it still exhibits wall-following. Smithers (1994) concludes that we must move beyond characterizing sensors as measurement devices. Sensor signals do not ‘encode information’ specifying states of a robot in its environment. What they do is “vary in some way that depends upon the dynamics of the robot–environment interaction.”

This novel view of sensory-motor functioning may help to take our explanations of cognition and action beyond the limiting directionality that was discussed above in relation to subject–environment dualism. At both synchronic and diachronic levels of analysis, a new viewpoint becomes possible. At the synchronic level, input–output notions of encoding, recovering or picking up invariant information are replaced by concern with an ongoing process of sensory-motor covariation. At the diachronic level, the traditional view sees a system that is reactive with respect to its environment becoming able to an-

ticipate what is going to happen through internal representations. This is questioned by the view that all interaction actually takes place in a dynamic interactive present, never in the past or future (Smithers, 1995), suggesting a new question: How can a system's dynamics change to take account of past history in a way that enables it to extend its dynamic interactive present and to generate the performance(s) that we associate with anticipation of the future?

This kind of rethink may more readily support genuinely enactive, mutual notions of perceptual organization, such as Koenderink's (1980) right-sounding yet mechanistically opaque claim that "you do not 'extract' what is already there: what is there depends on me.....I do not become attuned to things: the things are what they are because I am what I am."

References

- [1] Beer, R.D. (1995) A dynamical systems perspective on agent–environment interaction. *Artificial Intelligence*, 72, 173-215.
- [2] Brooks, R. (1991). Intelligence without reasoning. *Proceedings of the Twelfth International Joint Conference on Artificial Intelligence*.
- [3] Cliff, D., Harvey, I. and Husbands, P. (1994) General visual robot controller networks via artificial vision. University of Sussex, Cognitive Science Research Paper, Serial No.318.
- [4] Costall, A. (1995) Socializing affordances. *Theory and Psychology*, 5, 467-482.
- [5] Dennis, W. (1960) Causes of retardation among institutional children: Iran. *Journal of Genetic Psychology*, 96, 47-59.
- [6] Edelman, G.M. (1992) *Bright Air, Brilliant Fire: On the Matter of the Mind*. Allen Lane: Penguin Press.
- [7] Hess, E.H. (1956) Space perception in the chick. *Scientific American*, 195, (July), 71-80.
- [8] Husbands, P., Harvey, I. & Cliff, D. (1995) Circle in the round: State space attractors for evolved sight robots. *Robotics and Autonomous Systems*, 15, 83-106.
- [9] Keil, F.C. (1981) Constraints on knowledge and cognitive development. *Psychological Review*, 88, 197-227.
- [10] Koenderink, J.J. (1980) Why argue about direct perception? *Behavioral and Brain Sciences*, 3, 390-391.
- [11] Maturana, H. & Varela, F.J. (1988) *The Tree of Knowledge*. Boston & London: Shambhala.

- [12] Piaget, J. (1972) Intellectual evolution from adolescence to adulthood. *Human Development*, 15, 1-12.
- [13] Rader, N., Bausano, M. & Richards, M.E. (1986) On the nature of the visual cliff avoidance response in infants. *Child Development*, 51, 61-68.
- [14] Reeke, G.N., Finkel, L.H., Sporns, O. and Edelman, G.M. (1990) Synthetic neural modeling: A multilevel approach to the analysis of brain complexity. In G.M. Edelman, W.E. Gall & W.M. Cowan (Eds.) *Signal and Sense: Local and Global Order in Perceptual Maps*. Wiley-Liss.
- [15] Runeson, S. (1977) On the possibility of ‘smart’ perceptual mechanisms. *Scandinavian Journal of Psychology*, 18, 172-9.
- [16] Rutkowska, J.C. (1991). Looking for ‘constraints’ in infants’ perceptual-cognitive development. *Mind and Language*, 6, 215-238.
- [17] Rutkowska, J.C. (1993) *The Computational Infant: Looking for Developmental Cognitive Science*. Hemel Hempstead: Harvester Wheatsheaf.
- [18] Rutkowska, J.C. (1994a) Scaling up sensorimotor systems: Constraints from human infancy. *Adaptive Behavior*, 2, 349-373.
- [19] Rutkowska, J.C. (1994b) Unpacking the R-words. *Proceedings of the Third International Workshop on Artificial Intelligence and Artificial Life: Dynamics and Representation in Adaptive Behaviour and Cognition*. San Sebastian: EHU/UPV.
- [20] Rutkowska, J.C. (1995) Can development be designed? What we may learn from the Cog project. In: F. Moran, A. Moreno, P. Chacon and J.J. Merelo (eds.) *Advances in Artificial Life: Proceedings of the Third European Conference on Artificial Life*. LNAI/LNCS Series Number 696. Berlin, Heidelberg: Springer Verlag.
- [21] Scheier, C. & Pfeiffer, R. (1995) Classification as sensorimotor co-ordination: A case study on autonomous agents. In: F. Moran, A. Moreno, P. Chacon & J.J. Merelo (Eds.) *Advances in Artificial Life: Proceedings of the Third European Conference on Artificial Life*. LNAI/LNCS Series Number 696. Berlin, Heidelberg: Springer Verlag.
- [22] Simon, H.A. (1962) An information-processing theory of human development. In: W. Kessen & C. Kuhlman (eds.) *Thought in the Young Child*. Chicago: University of Chicago Press.
- [23] Smithers, T. (1994) On why better robots make it harder. In D. Cliff, P. Husbands, J.-A. Meyer & S.W. Wilson (Eds.) *From Animals to Animats 3*. MIT Press/Bradford Books.
- [24] Smithers, T. (1995) Are autonomous agents information-processing systems? In L. Steels & R.A. Brooks *The Artificial Life Route to ‘Artificial Intelligence’: Building Situated Embodied Agents*. Hove: Lawrence Erlbaum.

- [25] Thelen, E. & Smith, L.B. (1994) *A Dynamic Systems Approach to the Development of Cognition and Action*. Cambridge, Mass.: Bradford/MIT Press.
- [26] Varela, F.J. (1988). *Cognitive Science: A Cartography of Current Ideas*. Author's unpublished translation of F.J. Varela (1989). *Connaitre – Les Sciences Cognitives: Tendances et Perspectives*. Paris: Editions du Seuil.