

Computation, Dynamics and Sensory-Motor Development*

Julie C. Rutkowska
Cognitive & Computing

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, .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

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

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

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-representational 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 pregiven 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 pregiven 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

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.
- **Backing to evolution?** This framework places its emphasis on individual history, but

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 dis-

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

anticipate 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

