

Can Development Be Designed? What We May Learn from the Cog Project*

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

email: julier@cogs.susx.ac.uk
phone: +44 273 606755 x2403
fax: +44 273 671320

Abstract

Neither 'design' nor 'evolutionary' approaches to building behavior-based robots feature a role for development in the genesis of behavioral organization. However, the new Cog Project aims to build a humanoid robot that will display behavioral abilities observed in human infants; and proposes making use of ideas from evolution and developmental psychology in its design. This paper offers a provisional evaluation of this work from a developmental perspective, to show how developmental study may offer not only a source of phenomena for modelling but also a method that contributes to our understanding of how self-organization works. The design methodology that underlies Cog confronts problems with selection and interpretation of component behaviors, and how these may be better understood through appropriate developmental study is illustrated. Principles that underlie the design of Cog are shown to exhibit interesting convergences with infant mechanisms, based on the significance of emergent functionality and the action- as opposed to representation-based nature of both initial and outcome mechanisms. However, analysis of infants yields a more constructive view of ability, associated with different assumptions about the subject's relationship with the environment.

1 Routes to Understanding Autonomous Agents

Artificial Intelligence's new behavior-based robotics is unified by commitment to understanding intelligent systems in terms of specifics of their physical embodiment, their sensorimotor coupling with the environment, and the organizational possibilities of the situatedness to which these properties give rise. There is less agreement as to whether

*Submitted to the Third European Conference on Artificial Life. Granada, Spain, 4-6 June, 1995.

a historical process must be a necessary component of the construction of a system that is to become capable of survival in our normal environment. Engineering methods are at the heart of a ‘design’ approach to building robots, attempting to pre-specify component behaviors that are required and the mechanisms through which they can be implemented. Brooks’s insect-like Creatures (1986, 1990, 1991a & b), based on a subsumption architecture with layered control, provide elegant and successful examples of this strategy. Exponents of evolutionary robotics see this kind of hand-design as simply too hard to be feasible at any but a toy scale, however ingenious the experimenter. Instead, processes inspired by evolution are exploited to automate design. For example, notions like mutation, recombination and selection have been employed to evolve sensorimotor controllers in recurrent dynamical artificial neural networks by repeatedly evaluating and ‘breeding’ sets of (initially randomly generated) networks, thus arriving at a maximally adaptive genotype-like structure (Cliff, Harvey & Husbands, 1992, 1994; Harvey, Husbands & Cliff, 1992, 1994).

The overall role that a historical dimension plays in these approaches to autonomous organization is more complex than it may initially seem, at least as far as evolution is concerned. Superficially, they might appear to correspond to the distinction between “implementation by design” and “implementation by evolutionary strategies” that Varela (1988) identifies as one criterion of a theoretical shift from traditional symbol-manipulating cognitivism, focussing on the subject’s representation-based activities in a ‘pre-given’ world, to an enactive framework, in which the subject’s activities serve to construct a domain of interactions through which a world is ‘brought forth’.

On closer inspection, the gap is smaller than it seems. For example, both approaches purport to reject traditional representationalist assumptions (in favour of pragmatic ‘wiring’ considerations in the former case, and dynamical systems analysis in the latter). Furthermore, Brooks’s Creatures are deliberately engineered incrementally, while current evolutionary robotics pre-specifies behavioral outcomes insofar as selection is achieved through evaluation procedures that depend on objective, task-oriented fitness functions. Both approaches differ from autopoietic notions, which focus on how systems generate and maintain their own organization (i.e. are ‘self-producing’) and characterize evolution in terms of ‘natural drift’ rather than progressive adaptation to the environment (Maturana & Varela, 1988; Varela, 1988, 1993).

Neither approach features an explicit role for the individual historical dimension, development, in the genesis of behavioral organization. However, the ambitious new Cog Project aims to build a humanoid robot that will display behavioral abilities observed in the first couple of years of human life (Brooks & Stein, 1993; Brooks, 1994). Cog is not intended to be a model of human development. Nevertheless, it aims at biological relevance (Brooks, 1994); and proposes making use of ideas from both evolution and developmental psychology (Brooks & Stein, 1993). This paper aims to evaluate these proposals from a developmental perspective. Its emphasis is on going beyond the use of developmental study as a source of phenomena for modelling, to consider how it can provide a method that contributes to our understanding of self-organization.

2 From Creatures to Cog

The Cog Project is at an early stage, so any assessment must also be of a provisional nature. It aims to extend the methodology of generating complex seeming abilities through hard-wired networks of simple sensorimotor coordinations, each capable of engaging in independent interaction with the environment. As far as possible, such ‘task-achieving behaviors’ are to be added incrementally, with each of these layers being tested and debugged before attempting to build in another, continuing an analogy to evolution (Brooks & Stein, 1994, p.7). The initial schema for Cog includes far more layers than have so far been built into any Creature, offering an increased challenge for understanding behavioral coherence (Brooks, 1994). These layers begin with apparently basic abilities such as visual following and sound localization, and cumulate in complex ones such as generic object recognition and protolanguage. For example, layers most closely related to the ‘development’ of prehension are: body stability, leaning, resting; bring hands to midline; own hand tracking; hand linking; batting static objects; grasping & transfer; and body & arm reaching.

Cog “will get a continuous large and rich stream of input data of which it must make sense, relating it to past experiences and future possibilities in the world” (Brooks & Stein, 1993, p.6), but it appears to follow its Creature predecessors in lacking a clearly defined role for endogenous organizing mechanisms. Attaining novel abilities by building in additional layers is the consistent focus, in preference to more dynamic notions of designing the system so that novel abilities might ‘come for free’ as far as the design process is concerned, emerging from what happens when implemented layers are allowed to ‘run’ in the environment. For example, when it is assumed that functional behavior layers for correlating hearing and vision should serve as a usable basis for discrimination between ‘interesting’ events and background noise, this is ‘use’ more by the human designer than autonomously by Cog as part of an internal process of self-organization.

2.1 Cog, Evolution and Development

Basing layered control on an analogy with evolutionary development (Brooks & Stein, 1993) could be seen as implying a view of the evolution-development relationship based on the notion of terminal addition. The implications of this view are illustrated by a model that views Piagetian stages of intellectual development, from sensorimotor abilities to abstract thought, in terms of evolutionary selection pressures operating on genotypes that make possible specific behavioral adaptations (Parker, 1985; Parker & Gibson, 1979). Cross-species observations of relative attainments on developmental stages of ability are used as evidence that ontogeny recapitulates phylogeny, based on a series of terminal additions of new structures or stages through evolution. The fundamental argument is that “more intelligent species achieve their greater intelligence not by altering early developmental processes, but by adding later stages of intelligence to the end of the developmental cycle” (Gibson, 1981, p.52).

This view contrasts strongly with the epigenetic position favoured by Piaget (e.g.

1971), the origins of whose developmental stages it purports to explain, insofar as he considers stages as evidence for levels of knowledge that are neither additive nor genetically predetermined but the product of developmental processes that operate from the very outset. While Piagetian stages and developmental processes are contentious, there are good grounds for sharing his dissatisfaction with proposals for genetically pre-specified additive/sequential behavioral outcomes that arise from this way of using orthodox neo-Darwinian theory to frame an account of development. Attributing a privileged role to genes in the determination of development, commonly enshrined in the ‘genetic program’ notion, has appropriately been criticized as denying the very development that it seeks to explain (Oyama, 1985). Genes can be thought of as inputting certain parameter values into a developmental process involving a system of multiple variables and relations, but they do not define the organizing principles of that process. Those depend on the dynamics of the developing system as a whole (Goodwin, 1993). Observations of human acquisition of everyday, apparently universal sensorimotor abilities suggest processes capable of flexible outcomes that strain the notion of genetic predetermination. For example, infants who are raised with sparse adult interaction may not walk, even by around 3 years of age. Instead, they acquire ‘scooting’ (sitting while using arms to pull the body along), a behavior whose form is surely not pre-programmed (Dennis, 1960).

Brooks and Stein (1993) acknowledge that the analogy between layered control and evolution is ‘simplistic and crude’, and seem unlikely to wish to characterize development in the (over)simple additive, genetically bound terms sketched above. Their current models, however, appear compatible with that direction. The evolutionary approach to constructing robots is sensitive to the fact that its current models locate the form of individual performance too exclusively ‘in the genes’, and both approaches agree on serious reservations about viewing real evolution as an orthodox process of optimisation with contemporary animals seen as solutions to problems posed in their species’ distant evolutionary past (Brooks, 1994; Cliff, Harvey & Husbands, 1994). However, just how new sensorimotor coordinations emerge in development, if not through an essentially additive sequence of gene-behavior mappings, remains an open question. Some developmental ideas about how transactions between phylogenetically determined initial mechanisms and the environment may guide ontogenetic change are sketched in the following section of this paper.

2.2 Using Developmental Observations

Constructing Cog by design entails pre-selection of behaviors into which its abilities will be decomposed. Its planned layers thus embody an implicit developmental theory, to the extent that they highlight a restricted range of the behaviors that have been studied by developmental psychologists, and provisionally order them so that earlier layers are expected to aid the implementation and operation of later ones.

Brooks’s work emphasises the difficulty of achieving an effective behavioral decomposition of abilities, and how we may frequently be misled into thinking that our

observers' discriminations map straightforwardly onto demarcations in our subjects' mechanisms. Certainly, the selection and interpretation of behaviors for such a design plan raise closely related problems. On the one hand, behaviors that are necessary to the developmental sequence may be missed. In the case of prehension, for example, arm raising behavior is commonly found after batting objects is observed but before top-level reaching appears; while the infant intently fixates the object, the arm and hand are raised in its direction and held at the horizontal, often with signs of considerable effort. This behavior is not included in the Cog schema, although it may play a more significant role in the emergence of reaching and grasping than, say, hand linking. Equally, of course, it may not. Deciding which is the case depends on interpretation of the behavior concerned; that requires a hypothesis, or at least a hunch, as to what mechanisms are involved, which clearly affects plausible implementation strategies.

The difficulty of such interpretation as far as individual behaviors are concerned can be illustrated easily, without considering potential Cog layers that have controversial psychological connotations (e.g. that dedicated to 'multiple drafts emergence' associated with Dennett's ideas about consciousness). What is the significance, for example, of the readily observable behavior of bringing hands to the midline? There is no developmental consensus as to why infants exhibit midline activity (Rutkowska, 1994b). Possible interpretations range from initially out-of-sequence fine motor movements that will eventually be used for manipulating grasped objects; to stress reduction in the case of hands brought to the mouth; and a side-effect of the mechanics of failed early reaching attempts. Likewise, is batting at objects a form of ballistic reaching, superseded by reaching with visual feedback, or is it an attempt to palpate an extended surface that is perceived as too large for grasping?

A prime source of such difficulties is methodological. Much mainstream developmental psychology does not itself employ a particularly developmental method. It tends to concentrate on relatively isolated behaviors; is preoccupied with when in the subject's chronology those behaviors appear; and with judgments of 'success/failure' that treat behavior as a mere criterion of some other, supposedly underlying ability. And it generally works backwards from possibly erroneous assumptions about outcome behaviors to processes of acquisition (Rutkowska, 1993). Faced with the question 'How do you get from **A** to **B**?', it concentrates too exclusively on the nature of **B**. An essential change of focus is needed from initially asking 'Can this system do **B**?' to the more fundamental question 'What is this system doing?' This change of orientation is facilitated by employing a more genuinely developmental method.

The strategy involves three things that enrich, and often alter, our understanding of the nature of **B**:

- Taking behavior seriously by looking at the patterns through which it achieves succeed/fail outcomes in terms of the observer's criteria.
- Working forwards by taking seriously the idea that you can't get from **A** to **B** unless you start from a good idea of what **A** is.

- Observing changing behaviors in a domain of activity, using the relative position of a behavior within a sequence to constrain its interpretation.

3 Behavioral Interpretation Through Development

The developmental strategy can be illustrated by looking at infants' changing performance on a simple visual tracking task that presents them with a moving object, part of whose path is hidden by an occluder (Rutkowska, 1993, 1994a & c). Their looking behavior is generally assumed to index knowledge of the object and of its motion ('success' = look to exit as/before the object reappears; 'failure' = look elsewhere). Even very young infants will sometimes succeed in 'anticipating' the object's emergence from behind the occluder in an operational sense, by looking at the exit side as or before the object comes back into view. Should we therefore conclude, depending on theoretical preference, that infants come equipped with visual procedures for solving the problem of object search or 'believe' that objects continue to exist while out of sight? Considering the details of this behavior in the context of others displayed by 3-, 6- and 9-month-old infants makes such interpretations extremely implausible. Three aspects of the data are notable:

- The behavior pattern of fixations and head and eye movements that sometimes leads 3-month-olds to be looking at the object's reappearance point before it comes into view is quite different from the pattern through which 9-month-olds attain the same outcome. While 3-month-olds simply continue tracking as the object disappears from view, sometimes tracking as far as the reappearance point, 9-month-olds characteristically pause as the object disappears from view, then make a single head and eye movement to the reappearance side of the occluder, which they fixate until the object returns to view.
- Although 3-month-olds' continued tracking has the appearance of functional search for the disappeared object, its frequency declines rather than increasing with age. Nor is it simply replaced by a corresponding increase in the 'entry-exit' fixation pattern found in 9-month-olds, despite infants getting faster and faster at turning to refixate the reappearing object, from wherever they do happen to be looking, as it comes into peripheral vision. 6-month-old subjects exhibit less of either form of 'successful' anticipation than 3- or 9-month-olds, demonstrating the kind of U-curve that characterizes many instances of development.
- What does increase are behavior patterns involving attention to the object's disappearance point. The one most characteristic of 6-month-olds can be described as backtracking: as the object disappears, the infant continues tracking, but then turns head and eyes sharply back to fixate the object's disappearance point. This is a strange observation as far as attempts to interpret backtracking in isolation are concerned, since those generally assume the infant must have noticed some change in the reappearing object and be looking back to the disappearance point

where the original object was last seen. Here, however, a single object moving at constant speed is involved, and is generally still out of view when the infant turns back.

These and other aspects of the data suggest the observer-labelled tracking task is not initially a problem with the goal of ‘find the object’ as far as the infant’s viewpoint is concerned. 3-month-olds’ behavior is not wired up to search for objects that move and disappear, but they are initially equipped with a preadapted behavioral procedure for tracking visible object movement. Their continued tracking is no more than a failure to alter ongoing behavior when environmental circumstances change. They fail to do so because the recurrent visual pattern (kinetic occlusion) that marks the moving object’s disappearance has yet to become salient to them. It may be available preattentively, at the sensory process level, but has yet to be usable at the level of action through coordination with behavior. If this is the case, we would expect 3-month-olds to do one of two things when the object moves out of sight: nothing, i.e. look away or ‘lose interest’; or what they are already doing, i.e. continue tracking. These prove to be the behavior patterns that are most characteristic of that age. 6-month-olds’ backtracking can be seen as indicating the beginnings of attention to kinetic occlusion, which develops further with behavior patterns such as intently fixating the object’s disappearance point during the entire period that it is out of sight. They are not seeking a changed/missing object. Only the 9-month-olds’ coordination of attention to kinetic occlusion with turning to look to the opposite side of the occluder marks the beginnings of search from both the observer’s and the infant’s viewpoints.

Infants, then, are not initially trying to ‘do the task’ as an observer sees it. Their changing performance is more akin to task construction than task solution. Through repeated sensorimotor interactions with the environment, they come to construct the problem of search for missing objects through their experience of finding them. Even at this everyday level, development may be seen as an enactive process, in Varela’s (1988) terms, insofar as its processes appear to be directed at problem-definition rather than problem-solving.

4 Action-Based Task Construction

The broad issue of how novel abilities are constructed can be crudely divided into three questions, each of which allows additional evaluation of the architectural principles and design methodology of the Cog Project from a developmental perspective: What are initial mechanisms like?; What is the process through which they change like? And what do they change to?

4.1 Emergent Functionality in Scaffolding Development

A key point of rapprochement between design and developmental methodologies involves the notion of emergent functionality, through which complex abilities may result from the independent interaction of more basic components with the environment.

Emergent functionality is central to the Cog Project's attempt to maintain behavioral organization through layered control, and it may be developmentally advantageous in two ways, at least as far as the early stages of acquiring novel abilities are concerned (cf. Rutkowska, 1994a & c).

Firstly, emergent functionality could support an initial organization of independent sensorimotor coordinations, such as the visual following featured in the preceding section, that is neither a *tabula rasa* nor a blanket prewired solution to problems that will be encountered. This would offer preadaptation without rigid predetermination. Interactions between preadapted abilities of such a system and the environment in which it finds itself could enable it to 'tune in' sensorimotor coordinations, and sequences of such coordinations, that prove viable in the individual's experience. Novel coordinations (e.g. locomotion by scooting) would not be precluded in case of altered environmental conditions and/or properties of the subject (e.g. physical-motor disability).

Secondly, within the developmental process, the phenomenon of scaffolding can be viewed as a form of supervised learning in which emergence of function is temporarily engineered to establish the developmental space within which viable patterns of activity can be stabilized. Scaffolding, as originally viewed in social terms, marks the process through which more able humans manipulate the infant's transactions with the environment so as to foster novel abilities (e.g. Valsiner, 1987; Wood, Bruner & Ross, 1976).

The process begins with sensory and motor processes that are not coordinated by the infant but are set in alignment with the environment by adults. For example, if an infant's head is moved to look at someone leaving a room and simultaneously his/her hand is moved up and down, whatever the infant is doing, initially s/he is not waving goodbye. Key features are: customizing or simplifying the environment; reducing the number of degrees in the target task; directing attention by marking critical attributes; and enabling repeated experience of the end, outcome or goal of an activity that the infant would be unable to seek voluntarily. This sets up the possibility of serendipitous learning by the infant, that is of an accidental (i.e. unplanned) yet fortunate discovery of possibilities for effective action, in which the balance of behavioral control shifts from the environment to the subject.

The ubiquitous nature of such phenomena has been seen as evidence for all aspects of human development being socially and culturally guided, but adults may be exploiting and directing inbuilt processes that also operate in infant's spontaneous interactions with the environment. For example, in the previous section's account of the development of visual tracking, initial serial ordering of behaviors emerged from ongoing interaction with the environment; it was not governed by a goal or plan directed at finding the disappeared object. Spatio-temporal properties of the infant's interactions with the environment supported recurrent sequences of sensory and motor processes, most notably attention to kinetic occlusion followed by turning to refixate (and hence to experience 'finding') the reappearing object. In principle, such processes may share the main properties of social scaffolding, provided attention can be limited through

processing restrictions such as spatiotemporal constraints (for a relevant simulation of an attention mechanism in the context of sensorimotor learning see Foner & Maes, 1994).

The notion of scaffolding begins to provide a way out of problems faced by traditional AI's view of learning. This tends to see it in terms of adaptive change that enables a system to do a task better next time round; and which is unnaturally difficult unless the subject knows the goal in advance (Mitchell, 1983). Such assumptions make it difficult to see where novel abilities and goals might come from. It is notable that a robot system such as Darwin III, which is purported to exhibit self-organization in the absence of supervision and with unbiased internal connectivity in place of inbuilt sensorimotor structure, relies heavily on designer-coded 'value schemes' that evaluate the outcomes of its behavior (Reeke, Finkel, Sporns & Edelman, 1990). These intrinsic value schemes (e.g. getting visual stimuli onto the fovea = good; making contact with bumpy objects = bad; hand in region of foveated object = good) share non-trivial properties with traditional internal goals (and with externally specified fitness functions), hence encounter similar problems. As a means of ensuring recurrent experience of novel viable activities, coupling emergent functionality with scaffolding may offer a better characterization of constraints on infant's changing behavior.

4.2 Is Development Additive?

The kinds of behavioral developments that characterize the first year of life appear to involve more than straightforward addition of novel sensorimotor coordinations. In many domains of activity, there appears to be a move away from a reactive mode, which is essential to the basic operation of layered control in Brooks's robots, to increasingly anticipatory functioning and what might be called 'nascent plans'. In the case of visual search, for example, the infant develops from attention to kinetic occlusion plus turning to fixate the reappearing object to attending to kinetic occlusion then turning *in order* to re-fixate the reappearing object.

An explanation of the novel mechanisms that underlie such changes need not invoke qualitative change in the form of concepts or mentally represented goals controlling behavior. It will need to account for an extended time-scale of coupling between the subject's activities and the environment; and the changing functionality of sensorimotor components that is illustrated by eye and head movements initially associated with fixation coming to be used also for re-fixation. Traditional computational explanations could permit new and old 'programs' to invoke common lower-level movement primitives. It is, however, an interesting empirical question as to whether, and how, an intricately hardwired subsumption architecture could generate such phenomena.

4.3 What Do 'Internal World Models' Model?

Meyer & Guillot (1994) refer to anticipatory triples of the form '*if, in sensory circumstance 1, I do behavior B, I shall get sensory circumstance 2*' as 'internal world models'. While agreeing that such mechanisms underlie a form of nascent planning, it

is worth emphasizing that, to the extent that they ‘model’ anything, it is constraints on effective action rather than an external ‘world’ in which action takes place. Developmental psychology and cognitive science have become relatively fixated on the notion that model-like internal representations of the environment underlie intelligent functioning, a notion that is not endorsed here. Nor are representations featured in Brooks’s (1991b) foundational assumptions about intelligence, although it has been argued that his robots may in fact use internal representations and require them for further progress to be made (e.g. Clark & Toribio, 1994).

I doubt this conclusion as far as both infants and robots (whether humanoid or otherwise) are concerned. This is not because there is no interpretation of ‘representation’ that can be mapped onto their functioning, but because the notion of an internal representation attempts to demarcate one component of a complex subject-environment system, and to give it a privileged status in the genesis of organization. In doing so, it limits attempts to deepen understanding of how that organization is achieved.

Insofar as it works to establish selective correspondence(s) between subject and environment, action maps onto some perfectly good treatments of representation and the establishment of meaning, which are equally applicable to human infants, to Brooks’s Creatures and, by extrapolation, to Cog (Rutkowska, 1994a & c). What these action-based, process-oriented approaches share is a scale of analysis that spans sensory and motor processes and their functional coordination in the environment, unlike traditional preoccupations with representation as a substitution for the environment, which locates it firmly ‘in the head’. A typical direction is the situation semantics notion of ‘attunement to constraints’ (Israel, 1988), which allows infant sensory-motor coordinations, such as reaching towards or avoiding things, and the task-achieving behaviors of Brooks’s robots (e.g. 1991b) to be thought of as underlying human and robot subjects’ representation and understanding of constraints on acting in the world, and their ability to satisfy them. Also significant are Dretske’s (1988) notion of a ‘natural system of representation’, which can be applied to cases of infant action, as when very simple directionally selective elements acquire the function of indicating that something is approaching from the way they are ‘wired’ with avoidance behaviors; and to robots, as when a sonar pattern associated with free space acquires the function of indicating a place to visit when wired into a task-achieving behavior that successfully embeds the robot in its environment. Also applicable is Varela’s (1993) view of meaning emerging from processes that establish domains of interaction between a ‘self’ and its environment, for which the CNS’s sensory-, motor- and inter-neurons are only one specialist adaptation for achieving closure, a reflexive interlinking of subject and environment processes that supports construction of neurocognitive identity.

As far as such systems are concerned:

- Explicit internal models of an objective external environment are unnecessary to adaptive behavior, provided embedded sensory and motor processes are taken as the scale of analysis. The usefulness (efficiency) of sensory processes lies not in how exhaustively they enable the subject to model an object or an event but in how successfully they limit and allow for possibilities for action.

- No ‘bit’ of action mechanisms is ‘the’ internal representation. The capacity for successfully locking onto the environment and anticipating the consequences of activity within it is distributed across the operation of perceptual and behavioral processes. For example, infants’ knowledge of invisible object movement is embedded in the way the disappearance event becomes involved in determining future head and eye movements towards the reappearance point.
- The notion of ‘representation’, when viewed in terms of action-based mechanisms, does no explanatory work in the sense of being a (more or less localizable) functional component of action. Representation, whether by selective correspondence or by substitution, is one vantage point onto processes that are grounded in an action system spanning subject and environment. It makes equally little sense to consider this a central/internal phenomenon or an external one.

When it comes to considering the environmental contribution to the interaction of subject and environment, it seems that Brooks’s methodology does not fully follow through the implications of this systematic, action-based approach. Whereas traditional ‘representation’ was appropriately dismissed, its opposite number, ‘information’, is invoked in vestigial but non-trivial allusions to animals having sensors that “extract just the right information about the here and now around them” (Brooks, 1991a). Despite a subjective focus on embodiment and action, a notion of an objective environment containing information for action is brought in, albeit information that is selected with the subject’s action requirements in mind. This can raise inappropriate assumptions about the subject’s ‘access’ to the world.

In the context of challenging model-like representations, work in behavior-based robotics has often alluded approvingly to Gibson’s (1979) theory of direct perception and the notion of ‘affordances’ (invariant combinations of properties that specify what things are ‘for’ for a given subject). It is often unclear just how far this flirtation with ecological psychology’s methodological realism is intended to go. Buying wholeheartedly into the view that there is an organized world of objects and events, independent of the subject’s activity, and that this is the world that is perceived and known, risks simply substituting a discredited notion of internal representations about the world for equally dubious external information about the world as ‘the’ foundation for organization. This can only be counterproductive as far as establishing a genuinely systematic account of adaptive behavior is concerned (cf. Bersini, 1992). It can make for easier questions for robot design and for development (though they may not be appropriate questions). To understand a subject’s world, we do not need to ask how organization is constructed, only about what objects and events they are sensitive to, and about what information specifies those things. Likewise, development is not a constructive process but one of coming to perceive better through the ‘education of attention’ to things that are already ‘out there’.

If we reject the information pickup/recovery metaphors for perception, and assume the environment may not come ready populated with the categories of objects and events about which language-using scientists and philosophers routinely speak, there

may still be no need to be pessimistic about how we might talk about the environment. If the aim is to understand emergent organization within an environment, as opposed to what is ‘really out there’ beyond the subject and the subject’s activities, behavior-based robotics suggests some interesting possibilities.

One route comes from evolutionary robotics’ focus on the adaptation of simple robots in simple environments (e.g. Cliff, Husbands & Harvey, 1992). The power of this method lies in its potential for ‘reverse engineering’ to clarify what sensorimotor control structures the robot evolves for itself, in place of checking its success/failure in acquiring the categories of its creators. A further direction involves turning the issue of what sensors are doing away from the heavily entrenched notion that they provide environmental information (whether for direct action control or further processing). Dynamical systems theory may support a promising alternative if it can develop the notion that sensors are not measurement devices. It suggests that sensor signals need not ‘encode information’ specifying particular states of a robot in its environment, it is enough that they “vary in some way that depends upon the dynamics of the robot-environment interaction” (Smithers, 1994, p.70), delivering what might be considered sensorimotor invariants. This kind of re-think may provide an original route into genuinely enactive, mutual notions of organization.

5 Conclusion

Not only the successes but also the failures of the Cog Project will offer significant opportunities to clarify our thinking about development. As far as principles underlying the design of Cog are concerned, it starts out exhibiting some interesting convergences with the human infants whose behavioral abilities it aims to model, revolving around the significance of emergent functionality and the action- as opposed to representation-based nature of both initial and outcome mechanisms. As far as the design methodology is concerned, both selection and interpretation of Cog’s component behaviors pose difficulties, which might be eased by developmental data (of which there is less than one would like) that takes seriously behavior and the process of development. Along the lines of evolutionary processes enabling the construction of structures that may be too hard to achieve by hand design, exploiting developmental change may enable appropriate interpretation of behaviors and acquisition possibilities that are too hard to achieve purely through rational analysis. An ideal aim for Cog would be to reduce the design element in favour of carving out a greater role for an internally driven contribution to its own ‘development’ through interaction with the social and physical environment. This would offer a unique opportunity to clarify how both the outcome and the process of development are grounded in effective action.

References

- [1] Bersini, H. (1992). Animat’s I. In F.J. Varela & P. Bourginé (Eds.), *Towards a*

Practice of Autonomous Systems: Proceedings of the First European Conference on Artificial Life. Cambridge, MA: MIT Press/Bradford Books.

- [2] Brooks, R. (1986) A robust layered control system for a mobile robot. *IEEE Journal of Robotics and Automation*, **RA 2**, 14-23.
- [3] Brooks, R. (1990) Elephants don't play chess. In P. Maes (ed.) *Designing Autonomous Agents*. Bradford/M.I.T. Press.
- [4] Brooks, R.A. (1991a) Intelligence without reasoning. *Proceedings of the Twelfth International Joint Conference on Artificial Intelligence*.
- [5] Brooks, R. (1991b) Intelligence without representation. *Artificial Intelligence*, **47**, 139-160.
- [6] Brooks, R.A. (1994) Coherent behavior from many adaptive processes. In D. Cliff, P. Husbands, J.-A. Meyer & S.W. Wilson (Eds.) *Animals to Animats 3: Proceedings of the Third International Conference on Simulation of Adaptive Behavior*. Cambridge, MA: MIT Press/Bradford Books.
- [7] Brooks, R.A. and Stein, A. (1993) Building brains for bodies. MIT AI Laboratory Memo No. 1439.
- [8] Clark, A. and Toribo, J. (1994) Doing without representing? University of Sussex, Cognitive Science Research Paper, Serial No. CSRP 310.
- [9] Cliff, D., Harvey, I. and Husbands, P. (1992) Incremental evolution of neural network architectures for adaptive behavior. University of Sussex, Cognitive Science Research Paper No.256.
- [10] Cliff, D., Harvey, I. and Husbands, P. (1994) General visual robot controller networks via artificial vision. University of Sussex, Cognitive Science Research Paper No.318.
- [11] Dennis, W. (1960) Causes of retardation among institutional children: Iran. *Journal of Genetic Psychology*, **56** 77-86.
- [12] Dretske, F.I. (1988) *Explaining Behavior*. Cambridge, MA: MIT Press/Bradford Books.
- [13] Foner L. and Maes, P. (1994) Paying attention to what's important: Using focus of attention to improve unsupervised learning. In D. Cliff, P. Husbands, J.-A. Meyer & S.W. Wilson (Eds.) *Animals to Animats 3: Proceedings of the Third International Conference on Simulation of Adaptive Behavior*. Cambridge, MA: MIT Press/Bradford Books.
- [14] Gibson, J.J. (1979) *The Ecological Approach to Visual Perception*. Boston MA: Houghton-Mifflin.

- [15] Gibson, K.R. (1981) Comparative neuro-ontogeny. In G. Butterworth (ed.) *Infancy and Epistemology*. Brighton: Harvester.
- [16] Goodwin, B. (1993) Development as a robust natural process. In W. Stein & F.J. Varela (eds.) *Thinking About Biology*. Reading, MA: Addison-Wesley.
- [17] Harvey, I., Husbands, P. and Cliff, D. (1994) Issues in evolutionary robotics. University of Sussex, Cognitive Science Research Paper No. 219.
- [18] Harvey, I., Husbands, P. and Cliff, D. (1994) Seeing the light: Artificial evolution, real vision. In D. Cliff, P. Husbands, J.-A. Meyer & S.W. Wilson (Eds.) *Animals to Animats 3: Proceedings of the Third International Conference on Simulation of Adaptive Behavior*. Cambridge, MA: MIT Press/Bradford Books.
- [19] Israel, D. (1988) Bogdan on Information. *Mind and Language*, **3** 123-140.
- [20] Maturana, H. and Varela, F.J. (1988) *The Tree of Knowledge*. Boston & London: Shambhala.
- [21] Meyer, J.-A. & Guillot (1994) From SAB90 to SAB94: Four years of animat research. In D. Cliff, P. Husbands, J.-A. Meyer & S.W. Wilson (Eds.) *Animals to Animats 3: Proceedings of the Third International Conference on Simulation of Adaptive Behavior*. Cambridge, MA: MIT Press/Bradford Books.
- [22] Mitchell, T.M. (1983) Learning and Problem Solving. *Proceedings of the Eighth International Joint Conference on Artificial Intelligence*.
- [23] Oyama, S. (1985) *The Ontogeny of Information*. Cambridge: Cambridge University Press.
- [24] Parker, S.T. (1985) Higher intelligence as an adaptation for social and technological strategies in early Homo Sapiens. In G. Butterworth, J.C. Rutkowska & M. Scaife (eds.) *Evolution and Developmental Psychology*. New York: St. Martin's Press
- [25] Parker, S.T. and Gibson, K.R. (1979) A developmental model for the evolution of language and intelligence in early hominids. *Behavioral and Brain Sciences*, **2** 367-407.
- [26] Piaget, J. (1971) *Biology and Knowledge*. Edinburgh: Edinburgh University Press.
- [27] 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*. New York: Wiley-Liss.

- [28] Rutkowska, J.C. (1994a) Emergent functionality in human infants. In D. Cliff, P. Husbands, J.-A. Meyer & S.W. Wilson (Eds.) *Animals to Animats 3: Proceedings of the Third International Conference on Simulation of Adaptive Behavior*. Cambridge, MA: MIT Press/Bradford Books.
- [29] Rutkowska, J.C. (1994b) Prehension intention from 12 to 22 weeks. Presented at the IXth International Conference on Infant Studies. Paris, 2-5 June.
- [30] Rutkowska, J.C. (1994c). Scaling up sensorimotor systems: Constraints from human infancy. *Adaptive Behavior*, 2, 349-373.
- [31] Smithers, T. (1994) On why better robots make it harder. In D. Cliff, P. Husbands, J.-A. Meyer and S.W. Wilson (eds.) *From Animals to Animats 3*. Cambridge, MA: MIT Press/Bradford Books.
- [32] Valsiner, J. (1987) *Culture and the Development of Children's Action*. Chichester: Wiley.
- [33] 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.
- [34] Varela, F.J. (1993). Organism: A meshwork of selfless selves. Second European Conference on Artificial Life. Brussels, 24-26 May.
- [35] Wood, D., Bruner, J.S. and Ross, G. (1976) The role of tutoring in problem-solving. *Journal of Child Psychology and Psychiatry*, 17 89-100.