Chapter
18

Are Humans Mobot Minded? Some Implications of Embodied Cognitive Science for Cognitive Psychology¹

Stephen R. Hill

A mobot is a mobile robot. In the past decade or so a group of breakaway researchers have built new kinds of mobots that incorporate non-traditional control architectures and make heavy use of bodily and environmental resources in the production of adaptive behaviour (Pfeifer, Iida, & Bongard, 2005). *Situated robotics*, as this new perspective is often known (Mataric, 2003), forms a central plank of a new approach to studying cognition called *embodied cognitive science* (ECS; Anderson, 2003; Clark, 1997; Keijzer, 2001; Pfeifer & Scheier, 1999; Sporns, 2003; Wilson, 2002). The primary unit of analysis for much ECS research is the *behavioural system*, which consists of dynamically coupled neural, bodily, and environmental components (Chiel & Beer, 1997; Keijzer, 2001; see also Wilson, 2002, for criticism of this idea). ECS theorists are particularly concerned with the phenomenon of *behavioural emergence* (Hendriks-Jansen, 1996), the idea that behaviour is not entirely controlled by an 'in-the-head' control system but rather arises 'from the ongoing interaction between [an agent's] brain, its body, and its environment' (Beer, 2003, p. 210).

One key question raised by this novel approach is whether natural cognisers (people and other animals) might be better understood as mobot-like *embodied* and *situated* systems rather than the reflective, world-modelling systems beloved of classical cognitive science. Both embodiment (the possession of a physical body made up of components with their own dynamics and transduction features) and situatedness (the fact that agents rely on and make use of properties of their environments in the production of behaviour) provide strong constraints on the structure of the internal architecture ('minds') of situated robots. Embodiment constraints can be further decomposed into *integration constraints* and *morphological constraints*. The former arise from the demand that a complex

agent's 'brain' must combine multiple behaviour-producing subsystems into a coherent whole. The latter are a consequence of agents possessing non-neural body structures. The constraining effects of situatedness can also be examined in two ways. Many discussions of situatedness focus on *environmental constraints* — ways in which an agent's environment is used to simplify the computational demands of, and thereby the structure of the internal mechanisms that underpin, various tasks. An equally important aspect of situatedness is the concept of *interaction space*. This term refers to the ways in which important stimuli are temporally filtered so that cognitive abilities can robustly develop.

This chapter explores the idea that the design of the human mental architecture, like those of situated robots, is a product of integration, morphological, environmental, and interaction space constraints. Although much situated robotics research has focused on the generation of rather basic behaviours, rather than the kinds of complex cognitive abilities central to cognitive psychology, I want to suggest that we have much to learn from the synthetic methodology used in ECS.

Background: Mobile Robots and Cognitive Architecture

Artificial intelligence (AI) researchers have had a longstanding interest in creating robots that are self-sufficient and autonomous. However, by these standards, the early days of mobile robotics (the late 1960s and early 1970s) were not a resounding success (see Brooks, 1999). AI has attempted to circumvent the problems that gave rise to these failings in two ways: by either (a) focusing on simulating or building disembodied, function-specific systems known as expert systems, or (b) by simplifying the environments in which the classically-designed robots 'live'. In both cases, traditional assumptions about the structure and dynamics of internal processing systems remain largely unchanged.

In the mid-1980s the roboticist Rodney Brooks began to wonder whether the problem of building robust, self-sufficient robots could be due to traditional assumptions about the nature of internal mechanisms and processes. He set out to create cheap robots that could survive unaided in real environments filled with things like chairs, desks, doors, Coke cans, and graduates. On the face of it this was an incredibly ambitious undertaking considering the difficulties earlier robots had simply trying to cope with recognising and fetching a few geometrical shapes in uncluttered and simplified environments. It was amazing then that Brooks came up with a design that achieved many of his original objectives. The result was Allen, a simple wanderer of the MIT labs. 'He' was the first in a long line of increasingly sophisticated real-world MIT mobots (see Breazeal & Scassellati, 2002; Brooks, 1999).

The important point, for the purpose of this chapter, is that the innards of Brooks' robots bear almost no resemblance to the designs favoured by classical cognitive scientists (see also Harvey et al., 2005; Pfeifer et al., 2005; Pfeifer & Scheier, 1999). The reasons for this derive largely from the fact that Brooks' machines are both *embodied* and *situated* entities. Put simply, the task of building

complete, physically-realised robots, which can 'survive' in real environments, has constrained the design of the robots' computational control structures.

Layered Control Architectures and Integration Constraints

A basic assumption that classical AI shares with traditional cognitive psychology is the view that cognitive architectures are best understood as sequential information-processing systems made up of complex data structures that constitute models, plans, and goals (Palmer & Kimchi, 1986). The standard information-processing architecture involves a cascade of processing that usually includes sensing, modeling, planning, and acting stages. In this sequence, information from the environment is gradually filtered by a variety of algorithms and augmented with stored knowledge until motor commands for producing appropriate behaviours are computed. This sense-model-plan-act (SMPA) scheme also lies at the heart of much thinking in cognitive psychology. Its presence is obvious, for instance, in the sequence of detection, recognition, and response selection phases psychologists have used to understand the duration of mental processes since the time of Donders and Helmholtz.

Many of the researchers within the field of situated robotics have eschewed the SMPA schema and instead have built robots utilising *layered control architectures* (LCAs). Robots built using LCAs are, on the whole, better at 'surviving' in real-world environments than are traditional machines. The best known LCA is Brooks' *subsumption architecture* (for other architectures that adopt this principle of *parallel, loosely coupled processes* see Pfeifer & Scheier, 1999). LCAs typically omit the modelling and planning phases of processing (what Brooks refers to as the 'cognition box') and link sensing directly with action. Brooks argues that building a sufficiently useful model of the world is the most computationally expensive, and most problematic, part of the traditional AI scheme. He suggests that a system that can deal with changes in the environment in real-time will not necessarily require an internal model. As Brooks and his colleagues are fond of saying, such a system instead 'uses the world as its own model'.

Excising the 'cognition box' is one strategy for building more robust and adaptive robots. Another strategy is to do away with the need to build a system with a centralised control system. In traditional models of cognition, such as that offered by Marr's (1982) theory of vision, it is the job of the agent's mind/brain to construct detailed and centralised representations of the world which can be utilised in the execution of many possible actions. These representations are typically *detailed* (encoding information about many aspects of a stimulus) and *decoupled* (accessible by many response planning systems) (Clark, 1997, chap. 8; Sterelny, 2003). Brooks has argued that centralised control systems are unnecessary and unrealistic. Instead he favours a decentralised system decomposed into behaviour-producing layers. In this kind of system each layer receives minimal sensory input — just enough necessary for triggering a specific behavioural response (e.g., backing away from looming objects). Layers

Cognition and Language

Figure 1

Two examples of a layered control architecture: (A) Schematic of the subsumption architecture of the mobile robot Allen (based on Brooks, 1999); (B) The neural substrate underpinning rat defensive behaviour (based on Fig. 3 of Prescott et al., 1999).

then compete with each other for control of the robot's actions via a network of low bandwidth connections (see figure 1A for a simple example).

Subsumption architecture robots are built in an incremental fashion, somewhat reminiscent of the way nervous systems have evolved (see Prescott, Redgrave, & Gurney, 1999). Robot design begins by building a basic layer that utilises a narrow range of sensory information to produce a simple, reflex-like response (such as avoiding nearby objects). Once the basic layer has been perfected and debugged an additional layer is built atop the first which adds to the complexity of the robot's behaviour. It operates in parallel with the first layer and receives its own, specially tailored, sensory information, and produces its own specific motor responses (such as wandering in randomly selected directions). The

second layer has access to the activity of the lower layer, and can either inhibit output from the lower layer or, less frequently, replace the lower layer's output with its own 'data'. The lower layer, however, cannot affect anything happening in the higher layer. Additional layers are added on top of lower layers following the same incremental strategy until the robot's behaviour reaches a sufficient level of complexity.

An important implication of this design is that one cannot necessarily model the structure and dynamics of a higher layer mechanism without knowing the structure and dynamics of lower systems and how the higher layer interacts with them. In fact, one is likely to 'over engineer' a higher layer unless the constraints provided by lower layers are taken into account. LCA-controlled mobots do not store or construct models of their environments in order to plan how to act. Rather, their behaviour emerges from the coupling of the mobots' simple reflex subsystems and the unique structure of the local environment. The seemingly complex and goal-oriented behaviour of these machines is not explicitly pre-specified (encoded, programmed, or discoverable) in the mobot's innards. To an observer it may *appear* that a situated robot is operating on a unified conceptual model of a complex world, when, in fact, its internal machinery does not share or bring together information from its different sensor channels.

Layered Control Architectures in Natural Cognisers

Brooks (1999) and others have suggested that the LCA strategy was partly inspired by phylogenetic and neurobiological considerations, but that the primary reason for building mobots with layered control systems was because LCAs provide a better mechanism for producing robust adaptive behaviour. It is perhaps surprising, then, that some neurobiological researchers have begun to suggest that important aspects of vertebrate nervous systems can be usefully understood as LCAs (e.g., Quartz, 2002). Prescott et al. (1999) suggest that the idea that the nervous system of vertebrates is composed of semi-independent, behaviour-producing layers was foreshadowed in the writings of the nineteenth century neurologist John Hughlings Jackson (1884/1958). They describe research that suggests that dissociations consistent with the LCA profile have been shown to exist in a number of species (rats, cats, and fish). Ethologists have also claimed that this kind of machinery underpins much animal behaviour (e.g., Baerends, 1976; see also Hendriks-Jansen, 1996). They argue that many animals do not possess unified models of their environments or centralised goal-based planning systems. Instead behaviours are produced in response to sign stimuli — once one behaviour is executed the animal becomes sensitised to another class of stimuli that 'call' other behaviours. These chains of activity are organised in such a fashion as to produce large-scale patterns of adaptive behaviour. These ideas are echoed in Milner and Goodale's (1995) discussion of the anatomy and function of frogs' visuomotor systems:

The fact that the frog possesses this parallel set of independent visuomotor pathways does not fit well with the common view of a visual system dedicated to the construction of a

Cognition and Language

unified representation of the external world. Although the outputs of the different visuo-motor systems need to be coordinated, it would clearly be absurd to suppose that the different actions controlled by these networks are guided by a single visual representation of the world residing somewhere in the animal's brain. There is clearly no single representation or comparator to which all the animal's actions are referred (p. 11).

Prescott et al. (1999) have explicitly argued that the neural circuits that underpin rat defensive behaviour exhibit LCA-like features (see Figure 1B). These features include:

- **Distributed control:** Different layers make distinct contributions to the overall production of defensive behaviour.
- **Behavioural decomposition:** Evidence suggests that distinct behaviour sub-systems are underpinned by separate (dissociable) neural circuits. For instance, there exist separate circuits for 'controlling' backing-off defensive behaviour and forward-locomotion defensive behaviour.
- **Increasing levels of competence:** Higher layers in the rat defence system seems to mediate more complex responses to increasingly distant (temporally and spatially) threats.
- **Communication between layers by subsumption mechanisms:** Higher layers modulate the activity of lower layers by way of simple processes such as inhibition and substitution.
- **Limited sensor fusion and lack of centralised models:** At least some layers of the rat defence system use distinct 'sensor processing channels'. Furthermore, there is evidence that information from individual sensor channels is rarely fused to generate more complex centralised models of the local environment.

The LCA perspective may cast light on a number of theoretical models within cognitive neuroscience including the distinction between 'what/where' visuo-motor systems in humans and other mammals (e.g., Milner & Goodale, 1995) and Damasio and Damasio's (1994) convergence zone theory of the neural substrate of memory. In summary, there exists evidence that vertebrate brains can be usefully thought of as layered control systems.

Morphological Constraints

Situated robotics researchers have also discovered that the structure and dynamics of a robot's body, including the placement and operation of its sensors and actuators, strongly constrains the design of the robot's internal control system. The roboticist Tim Smithers (1995) argues that '[i]f a particular robot is not well fitted morphologically to its tasks and environment, no amount of clever programming will ever overcome its deficiencies' (p. 142). Smithers has found that there are times when a robot can be best 'reprogrammed' by altering its body shape rather than changing its internal program structure. Similarly, Pfeifer and Scheier (1999) suggest that computation can be 'traded' for morphology, and that good body designs can reduce computational complexity by an order of magnitude or more.

A hopping robot built by Raibert provides a nice example of the importance of morphology in behaviour production. This one-legged robot moves by bouncing on a single compressible leg. The robot's control structure controls the direction in which the leg points when it hits the ground. The movement commands are used to balance the robot so it does not fall over, and tip the whole machine so that it moves in a certain direction. However, the control system does not provide instructions for all of the movements exhibited by the robot. The robot's body provides bouncing and tipping over 'for free'. The control system modulates this intrinsic dynamic. Global behaviour, then, is an outcome of the coupling of the dynamics of the control system and the dynamics of the robot's body. Raibert and Hodgins (1993) summarise these ideas when they write:

We believe that *the mechanical system has a mind of its own*, governed by the physical structure and the laws of physics. Rather than issuing commands, the nervous system can only make 'suggestions' which are reconciled with the physics of the system and the task. (p. 350, italics added)

A similar claim has been made by the psychologist Esther Thelen and her colleagues (Thelen & Smith, 1994; Smith & Gasser, 2005) in their research with infants. Up until about 2 months of age infants exhibit a spontaneous stepping reflex when they are held upright. At around 2 months this pattern disappears. It returns between 8 and 10 months when the child is beginning to support itself on two legs. Traditionally it was thought that the disappearance of this reflex was due to maturational changes in the neural motor control system of the infant. However, in several ingenious experiments Thelen showed that these changes were caused by morphological changes in the infant. By the age of 2 months the weight of an infant's leg has reached the point where it simply defeats the natural spring-like action of the muscles. Thus, the neuromuscular system cannot generate the stepping pattern. Thelen demonstrated this by effectively reducing the weight of the infant's legs (by immersing the infant's legs in water or putting them on a treadmill). This finding implies that much of neural and bodily machinery needed for stepping is intact in 'non-stepping' infants. Thelen argues that many other changes in human behaviour can be accounted for in a similar way as the outcome of the coupling of multiple neural and non-neural body systems. Non-neural systems have also been implicated as important constraints in active vision (e.g., Findlay & Gilchrist, 2003) and auditory perception (Chiel & Beer, 1997).

Environmental Constraints

The actual dynamics of body structures depend to a large degree on the environment in which they are 'used'. For instance, in microgravity Raibert and Hodgins' robot would obviously not be able hop in a terrestrial fashion. Similarly, the behaviour of Thelen's infants changes when they are placed in different environments. Many situated roboticists have argued that a robot's body and internal control system cannot be designed unless one knows about the

layout of its 'species-typical' environment. The philosopher Andy Clark (1997; see also Rowlands, 1999) expands on the cognitive consequences of this idea by suggesting that '[i]n general, evolved creatures will neither store nor process information in costly ways when they can use the structure of the environment and their operations on it as a convenient stand-in for the information-processing operations concerned' (p. 46).

One can see this principle at work in the design of Polly, a mobot designed to show tourists around the MIT AI labs (see Pfeifer & Scheier, 1999). This ecological niche possessed a number of recurring features which simplified the design requirements for Polly's internal control structure. For instance, Polly needed to be able gauge the distance of various objects that lay in front of 'her'. This is a very complex problem. However, because Polly's niche consisted only of flat floors, distant objects always appeared higher on her camera sensors than closer objects. Thus, distance was a simple function of an object's Y coordinate. Polly's internal architecture was designed to exploit the unique physics of her local environment.

ECS researchers argue that the cognitive architectures of humans and other animals are also constrained by the structure and dynamics of their species-typical environments. Humans, of course, have colonised every habitat on the planet and have even begun to move beyond it. We seem to be amazingly adaptable in a wide range of environments. Despite this high level of adaptability, there are many ways in which our basic abilities work best in specific environments that include features found in our environment of evolutionary adaptedness such as illumination from above (for perceiving shapes using shading; Kleffner & Ramachandran, 1992) and natural light (for colour perception; Shepard, 1994). Both of these examples demonstrate constraints on basic perceptual abilities, but our cognitive abilities may also be designed to exploit features of our species-typical environments.

A simple example will help to illustrate this point. The visual memory literature suggests that people can remember a moderate number of features of a visuospatial layout. One might think, based on previous laboratory research, that we would make the most of this ability and fill, or nearly fill, our working memory when engaging in visuospatial tasks. However, in a number of experiments Ballard and colleagues (Ballard, Hayhoe, & Pelz, 1995; see also Ballard, Hayhoe, Pook, & Rajesh, 1997) showed that participants avoided memory-intensive activity whenever they could. The task in their experiments involved replicating a modelled pattern of coloured blocks. Ballard et al. showed that much of the time their participants often did not 'store' any more than one feature of the model (i.e., the colour or position of a single block) at a time. Indeed, it became apparent that people had no recollection of a previously 'stored' feature once a new one was 'encoded'. Simply put, people stored one feature at a time and promptly emptied their visual memory as soon as that feature had served its purpose within the larger task. Ballard et al. suggest that these results imply that we do not build complex internal models of our surroundings when we can use the environment as a kind of surrogate memory store (see also Gray & Fu, 2004).

Thus, there is reason to believe that humans, like mobots, exploit the structure of the environment when engaging in cognitive activity. The nature of the local environment provides important constraints on the structure and dynamics of the cogniser's internal cognitive mechanisms.

The Role of Interaction Spaces in the Development of Cognitive Mechanisms

Cognitive psychology often proceeds by attempting to model the cognitive abilities of *adult* humans. Although this is an admirable goal it ignores the important contribution that historical (evolutionary and developmental) analyses provide. ECS research supports the idea that a good understanding of developmental change can help to constrain our speculations about the nature of neurocognitive mechanisms that underpin various *mature* cognitive abilities. If we know something about the mechanisms and behaviour of earlier forms we can make better guesses about the structure and dynamics of internal cognitive mechanisms.

Verschure, Voegtlin, and Douglas (2003) have recently shown that early forms of a robot control system modulate the later forms by constraining the effective environment, or *interaction space* (Keijzer, 2001), to which the learning system is exposed. They constructed a simple, three-layered robot forager whose task was to learn to approach food targets and avoid obstacles. Verschure and colleagues found that performance was improved with the activation of a higher contextual layer, which enabled the robot to 'remember' extended temporal patterns of the associations between locations, obstacles, and targets. Upon investigation they found that this contextual layer influenced a lower adaptive learning layer, not by providing *internal* signals that improved the learning processes in the adaptive layer, but via an external feedback loop. That is, the contextual layer modulated the robot's activity by ensuring that the adaptive learning subsystem was provided with 'friendly' sensor input that enhanced the robot's ability to learn. Put simply, the contextual layer ensured that the robot's sensors were in the right place at the right time to make the most of the robot's limited learning capacity.

Does anything like this happen in animals? Verschure et al. (2003) think so: 'We expect that behavioural feedback would affect perceptual systems and those structures that readout from these systems, both during early development and beyond' (p. 623). They describe research that shows that the development of sensory areas (the auditory cortex, 'place' fields in the CA1 region of the hippocampus) might be influenced by behavioural feedback. In a more psychological vein several researchers have also suggested that language learning may be dependent on the structuring of interaction space by earlier forms (e.g., Deacon, 1997; Newport, 1990).

Elman (1993) has attempted to model these aspects of language acquisition using connectionist simulations. He found that in order to successfully train his networks to deal with complex sentential structures, it was necessary to constrain the stimulus array of sentence exemplars. In particular he found that two

constraining strategies were helpful: *phasing training* and *phasing memory*. The former involved gradually increasing the difficulty of the training sentences, while the latter involved endowing the infant network with a small 'immediate memory' and letting it gradually grow in capacity over time. Both of these strategies can be understood to have real-world counterparts: the transition from infant-directed speech to more adult speech, and the maturation of memory (Deacon, 1997).

The net effect of both strategies was to provide the infant network with a gradually expanding world of stimulation that enabled it to avoid getting stuck in local minima solutions from which recovery was impossible. What is illustrated here is the importance of the role of temporal constraints on the structure and dynamics of internal behaviour producing systems. Earlier forms (e.g., limited memory, poor visual acuity, restricted focal length) partially structure the kind of interaction space, to which an agent, is exposed enabling the development of later forms of behaviour that would not be possible had the earlier forms been more precocious (see also Smith & Gasser, 2005; Turkewitz & Kenny, 1985).

In sum this is another example of the, by now familiar, story that without a proper appreciation of the various factors that may constrain the structure and function of underlying neurocognitive mechanisms our attempts at cognitive modeling will miss the mark. So, to the list of layered control structure, morphological, and environmental factors we can add the constraining effects of developmental change.

Conclusions and Questions

I began by asking whether humans may be, in some important sense, mobot-minded — that is, whether the insights from the engineering of the internal control structures of situated robots could cast any light on the ways in which the cognitive architectures of humans, and other animals, are structured. Since embodied cognitive science's novel robotic designs have often been more effective in creating lifelike behaviour than those machines designed using classical architectures it might be the case that the useful engineering constraints that have been unearthed in situated robotics research might also be applicable to natural cognisers like us. I have tried to make a case for considering the possibility that our traditional cognitive models are inadequately constrained, and may therefore provide an unrealistic picture of the mechanisms that underpin cognition. I have attempted to show that, in at least some cases, there is evidence that the constraints explicitly acknowledged within the embodied approach are operative in humans and thus that they are potentially important for theory development in cognitive psychology.

There are, of course, many unresolved questions that may prove damaging to my case. Paramount among these is the question of whether we can scale up the lessons of situated robotics to the level of advanced cognitive abilities (Edelman, 2003). Modern adult humans are amazingly flexible cognitive creatures. We are not purely reactive, situation-driven beings. Indeed we seem to be able respond

in an almost infinite number of ways to environmental 'stimuli'. Thus, one might wonder whether the ECS framework has anything useful to offer a distinctly *human* cognitive psychology. It is relatively simple to come up with examples where each of the previously identified themes of ECS do not hold: regions of mammalian nervous systems do not exhibit layered architectures (Prescott et al., 1999); there are times when radical changes to a person's morphology do not result in significant changes in behaviour; often people find it easy to 'cognise' in a variety of environments; and sometimes early developmental forms seem to play no role in the structuring of the learning environment. These kinds of examples could be taken as evidence against the importance of ECS claims. There are many good reasons why we should not be too quick to make this theoretical leap (see, for example, Hill, 2000) but I will focus on just one: the fact that human beings are a species of animal that has evolved to cope with specific hostile environments (Sterelny, 2003) strongly suggests that our mental architecture is not fundamentally designed to enable off-line, rational pondering but rather to produce adaptive behaviour in response to (mainly local) environmental events. The flexibility we *do* possess is more than likely built on top of, and out of, previous structures that evolved as adaptations to earlier environments. Modern abilities may have arisen from *terminal additions* to the mechanisms of earlier, more restricted, abilities. For instance, Hermer-Vazquez, Moffet, and Munkholm (2001) suggest that the development of human spatial ability (from a simpler system that uses information about the macroscopic shape of the environment to a more flexible system that also utilises non-spatial information) may reflect the coming on-line of new layers of complexity. Similarly, it is possible that the increasing dominance of tertiary cortical areas in our species involves the subsumption of 'lower' cortical and sub-cortical systems by new control layers (Deacon, 1997). Curtis and D'Esposito (2003), echoing the claims of Hughlings Jackson (1884/1958), have recently suggested that the human prefrontal cortex, the *sine qua non* of pure mentality, should be understood as complex integrator of sensory and motor processes that lies 'at the apex of the motor hierarchy' (p. 421). Put simply, these examples suggest that human cognition may consist of an off-line mind built out of situated and embodied components.

If these claims about modern human cognition are anywhere near the mark, then the ideas broached in this chapter should have important implications for the ways in which we understand human cognition. We must face the possibility that the machinery that enables our flexibility rests upon, and depends upon, a reactive and situated neurocognitive bedrock.

Endnote

- 1 This chapter is derived from a longer manuscript (Hill, 2005) available from the author on request.

References

- Anderson, M. L. (2003). Embodied cognition: A field guide. *Artificial Intelligence*, 149, 91–130.

Cognition and Language

- Baerends, G. P. (1976). The functional organization of behaviour. *Animal Behaviour*, 24, 726–738.
- Ballard, D. H., Hayhoe, M. M., & Pelz, J. B. (1995). Memory representations in natural tasks. *Journal of Cognitive Neuroscience*, 7, 66–80.
- Ballard, D. H., Hayhoe, M. M., Pook, P. K., & Rajesh, P. N. R. (1997). Deictic codes for embodiment in cognition. *Behavioral and Brain Sciences*, 20, 723–767.
- Beer, R. D. (2003). The dynamics of active categorical perception in an evolved model agent. *Adaptive Behavior*, 11, 209–243.
- Breazeal, C., & Scassellati, B. (2002). Robots that imitate humans. *Trends in the Cognitive Sciences*, 6, 481–487.
- Brooks, R. A. (1999). *Cambrian intelligence: The early history of the new AI*. Cambridge, MA: MIT Press.
- Chiel, H. J., & Beer, R. D. (1997). The brain has a body: Adaptive behavior emerges from the interactions of nervous system, body and environment. *Trends in Neurosciences*, 20, 553–557.
- Clark, A. (1997). *Being there: Putting brain, body, and world together again*. Cambridge, MA: MIT Press.
- Curtis, C. E., & D'Esposito, M. (2003). Persistent activity in the prefrontal cortex during working memory. *Trends in the Cognitive Sciences*, 7, 415–423.
- Damasio, A. R., & Damasio, H. C. (1994). Cortical systems for retrieval of concrete knowledge: The convergence zone framework. In C. Koch & J. L. Davis (Eds.), *Large-scale neuronal theories of the brain* (pp. 61–74). Cambridge, MA: MIT Press.
- Deacon, T. W. (1997). *The symbolic species: The co-evolution of language and the brain*. New York: W. W. Norton.
- Edelman, S. (2003). But will it scale up? Not without representations. *Adaptive Behavior*, 11, 273–275.
- Elman, J. L. (1993). Learning and development in neural networks: The importance of starting small. *Cognition*, 48, 71–99.
- Findlay, J. M., & Gilchrist, I. D. (2003). *Active vision: The psychology of looking and seeing*. Oxford, England: Oxford University Press.
- Gray, W. D., & Fu, W.-T. (2004). Soft constraints in interactive behaviour: The case for ignoring perfect knowledge in-the-world for imperfect knowledge in-the-head. *Cognitive Science*, 28, 359–382.
- Harvey, I., Di Paolo, E., Tuci, E., Wood, R., & Quinn, M. (2005). Evolutionary robotics: A new scientific tool for studying cognition. *Artificial Life*, 11, 79–98.
- Hendriks-Jansen, H. (1996). *Catching ourselves in the act: Situated activity, interactive emergence, evolution, and human thought*. Cambridge, MA: MIT Press.
- Hermer-Vazquez, L., Moffet, A., & Munkholm, P. (2001). Language, space, and the development of cognitive flexibility in humans: The case of two spatial memory tasks. *Cognition*, 79, 263–299.
- Hill, S. R. (2000). *Where is cognition? Towards an embodied, situated, and distributed interactionist theory of cognitive activity*. Unpublished PhD dissertation, University of Canterbury, New Zealand.
- Hill, S. R. (2005). *Situated robotic systems as models of human cognitive architecture*. Unpublished manuscript.
- Jackson, J. H. (1958). Evolution and dissolution of the nervous system. In J. Taylor (Ed.), *Selected writings of John Hughlings Jackson* (Vol. 2, pp. 45–75). London: Staples Press. (Original work published 1884)
- Keijzer, F. A. (2001). *Representation and behavior*. Cambridge, MA: MIT Press.
- Kleffner, D. A., & Ramachandran, V. S. (1992). On the perception of shape from shading. *Perception and Psychophysics*, 52, 18–36.

- Marr, D. (1982). *Vision: A computational investigation into the human representation and processing of visual information*. New York: W. H. Freeman.
- Mataric, M. J. (2003). Situated robotics. In Nadel, L. (Ed.), *Encyclopedia of cognitive science*. (Vol. 4, pp. 25–30). London: Nature Publishing Group.
- Milner, A. D., & Goodale, M. A. (1995). *The visual brain in action*. Oxford: Oxford University Press.
- Newport, E. L. (1990). Maturational constraints on language learning. *Cognitive Science*, 14, 11–28.
- Palmer, S. E., & Kimchi, R. (1986). The information processing approach to cognition. In T. J. Knapp & L. C. Robertson (Eds.), *Approaches to cognition: Contrasts and controversies* (pp. 37–77). Hillsdale, NJ: Lawrence Erlbaum.
- Pfeifer, R., & Scheier, C. (1999). *Understanding intelligence*. Cambridge, MA: MIT Press.
- Pfeifer, R., Iida, F., & Bongard, J. (2005). New robotics: Design principles for intelligent systems. *Artificial Life*, 11, 99–120.
- Prescott, T. J., Redgrave, P., & Gurney, K. (1999). Layered control architectures in robots and vertebrates. *Adaptive Behavior*, 7, 99–127.
- Quartz, S. R. (2002). Toward a developmental evolutionary psychology: Genes, development, and the evolution of the human cognitive architecture. In S. J. Scher & F. Rauscher (Eds.), *Evolutionary psychology: Alternative approaches* (pp. 185–210). Boston: Kluwer.
- Raibert, M. H., & Hodgins, J. K. (1993). Legged robots. In R. D. Beer, R. E. Ritzmann, & T. McKenna (Eds.), *Biological neural networks in invertebrate neuroethology and robotics* (pp. 319–354). San Diego, CA: Academic Press.
- Rowlands, M. (1999). *The body in mind: Understanding cognitive processes*. Cambridge, England: Cambridge University Press.
- Shepard, R. N. (1994). Perceptual–cognitive universals as reflections of the world. *Psychonomic Bulletin & Review*, 1, 2–28.
- Smith, L. B., & Gasser, M. (2005). The development of embodied cognition: Six lessons from babies. *Artificial Life*, 11, 13–20.
- Smithers, T. (1995). Are autonomous agents information processing systems? In L. Steels & R. A. Brooks (Eds.), *The artificial life route to artificial intelligence: Building situated embodied agents* (pp. 123–162). Hillsdale, NJ: Erlbaum.
- Sporns, O. (2003). Embodied cognition. In M. A. Arbib (Ed.), *The handbook of brain theory and neural networks* (2nd ed., pp. 395–398). Cambridge, MA: MIT Press.
- Sterelny, K. (2003). *Thought in a hostile world: The evolution of human cognition*. Malden, MA: Blackwell.
- Thelen, E., & Smith, L. B. (1994). *A dynamic systems approach to the development of cognition and action*. Cambridge, MA: MIT Press.
- Turkewitz, G., & Kenny, P. A. (1985). The role of developmental limitations of sensory input on sensory/perceptual organization. *Developmental and Behavioral Pediatrics*, 6, 302–306.
- Verschure, P. F. M. J., Voegtlin, T., & Douglas, R. J. (2003). Environmentally mediated synergy between perception and behaviour in mobile robots. *Nature*, 425, 620–624.
- Wilson, M. (2002). Six views of embodied cognition. *Psychonomic Bulletin & Review*, 9, 625–636.