Chapter X

Stretching Detection: The Unique Selling Point of Embodied Cognitive Science

Stephen R. Hill

Massey University, Palmerston North, New Zealand

SUMMARY

By contrast with classical approaches to cognitive science, embodied cognitive science is characterised by claims that cognition is the product of coordinated systems of distributed and decentralised cognitive resources, and that it is grounded in survival-oriented architectures. I suggest that these ideas derive from a deeper change of perspective from viewing the cognizer as a representer to that of a detector/actor. Unlike some of the more abstract issues considered by other theorists, this change in perspective has practical implications for the ways in which psychologists study and understand cognition. In particular, it raises the important question of how we are able to access information about unperceivable environments using mental machinery that is not primarily designed for that purpose. A framework for understanding cognition as ‘time/space stretching’ is presented and three strategies for enabling that stretching are described.

This chapter examines the question of how we might understand what cognition is from an embodied cognitive perspective. In mainstream cognitive science, cognition is glossed as something like ‘the processing of information’ and this leads to an inevitable focus on explanations that are couched in terms of the computational transformation of representations. Rarely do cognitive scientists devote much time to asking *why* animals cognize and what is achieved when they do it – perhaps the answers seem obvious, but I shall argue that seriously considering these questions provides important theoretical food for thought. Over the last decade or so researchers in half a dozen disciplines have heralded the coming of an *embodied* cognitive science (ECS) which differs in significant ways from the classical orthodoxy (Anderson, 2003; Clark, 1997; Keijzer, 2001; Wheeler, 2005; R. Wilson, 2004). ECS *does* ask why animals cognize and the answers arrived at give us a very different sense of what cognition is. Specifically, I will argue that ECS tells us that cognition – all of it – involves some form of *active detection* rather than representation. Put crudely, cognition involves knowing about things by ‘seeing’ them (and doing things with them). In order to get a handle on this, possibly puzzling sounding, idea it will be helpful to quickly describe the core claims of ECS.

ECS ‘re-visions’ cognition as an embodied and embedded activity, which is to say that the focus shifts from viewing cognition as in-the-head ratiocination carried out by a computational brain, to viewing cognition as that bundle of processes and activities that enable animals (including we human animals) to make their way in their world (Christensen & Hooker, 2000; Clark, 1997).

According to ECS cognitive activity is the product of the coordination of a variety of resources distributed across the brain, the body, and the 'sociomaterial' world – tools, artefacts, built niches, and, not least, other cognizers (Kosslyn, 2006; Sutton, 2006; R. Wilson, 2007). Those who study distributed cognition often argue that there exists an unacknowledged explanatory symmetry that holds between internal and external resources in a cognitive system. More radically, some claim that non-neural, external elements actually constitute part of the mind (e.g., Clark & Chalmers, 1998). Advocates of both positions claim that extra-neural resources relieve the brains of cognizers of some of the burden of 'cognitive processing' by unloading it on to the environment. Thus, the role of the agent's brain is no longer considered to be that of a self-contained reasoner but rather that of a coordinator and utiliser of multiple informational resources. Indeed several types of cognitive activity have been shown to involve cycles of perception and action that rely heavily on environmental resources to do part of the 'computing' (e.g., Ballard et al., 1997). Moreover, even those activities that seem to involve classic 'off-line' cogitation have been reinterpreted as 'private simulations' of perception-action cycles (Hesslow, 2002; M. Wilson, 2001).

Relatedly it is claimed that, since cognizers have evolved and have physical bodies with physical requirements, all forms of cognition are likely grounded in the capacities for keeping that body functioning in, often dangerous, environments (e.g., Brooks, 1999). Thus it is not surprising that the basic cognition that enables this 'getting about in the world' has a perceptual, action-oriented, and 'emotional' character (Clark, 1997). Moreover, argue ECS researchers, these basic cognitive abilities provide the raw materials out of which more 'advanced' varieties of cognition – reasoning, reflective thought, science – are constructed (e.g., Barsalou, Breazeal, & Smith, 2007).

ECS advocates have spent the last ten years or so enumerating phenomena that illustrate these themes. The central contention of this chapter is that all of these ideas are underpinned by a more basic and telling idea: that cognition should really be thought of as detection; by which I mean, that aspect of animal life that involves establishing and maintaining contact with the world in order to inform the generation of adaptive behaviour. I will unpack this idea in the following way: First I will show how many of the concerns of cognitive science relate to the idea that agents need to know about things located at various spatial and temporal distances from their bodies and that this ability is what is usually meant by cognition. I then suggest that instead of making sense of this phenomenon using the notion of mental representation we are better served by thinking of cognition as involving attempts by a cognizer to 'stretch' time and space so that a variety of out-of-perceptual-range events and objects can play a role in the modulation of behaviour. 'Stretching' refers to the variety of ways in which detection mechanisms can be used to inform the cognizer about the environment. Obviously detection mechanisms can access the local environment, but, I argue, ECS also gives us clues as to how detection systems can be redeployed to 'get information' about things that are distant in time and/or space. This, possibly odd sounding, position might be helpful as a kind of perceptual switch that gets us thinking about cognition ECS-style, but more importantly it demands some attempt to understand how time/space stretching (TSS) is possible without falling back on traditional representational explanations. I close by considering several ways in which animals might be understood as detectors, even of things that are classically understood to be out-of-range.

COGNITION AS TIME/SPACE STRETCHING

Figure 1. A cognitive time/space cone. Time is represented on the vertical axis, space horizontally. The cognizer is located in the centre of ‘the Present’. The possible cognitive avenues for accessing different regions of the time/space cone are indicated by arrows and italicized descriptions.

The majority of animals are not ‘stuck’ in the immediate here and now. They possess abilities for modulating their behaviour with respect to elements of their worlds that they are not currently in physical contact with. Some phenomena, such as attention and perception, are primarily concerned with how animals manage to know about, and thereby coordinate their behaviour with respect to objects, events, and properties in the local (or ‘in range’) environment; that is, things with which the agent is in causal, but often not direct, physical contact. Other cognitive phenomena concern the ways in which animals modulate their behaviour with respect to things that exist in, what I call, distant (or ‘out-of-range’) environments — environments with which the agent is not in direct causal contact. For instance, memory research involves understanding the plethora of ways in which past events and experiences shape current functioning; planning, problem-solving, and reasoning often involve the modulation of behaviour with respect to anticipated, future circumstances; and language use often involves modulating behaviour by using information about distant environments. Unlike non-living things, people and other animals possess the power to sense the world about them, learn from previous experiences, make use of these ‘memories’, and, probably most importantly, anticipate the future (Suddendorf & Corballis, 2007). In a sense living systems are characterised by the fact that, to greater or lesser degrees, they ‘stretch’

themselves over spans of space and time incorporating and integrating the influences of distant things into their present activities.

This state of affairs can be visualized as a time/space cone (see Figure 1). In this characterization the cognizer (or more accurately, their physical body) is located in the centre of the immediate here and now, spatial distance is represented in the horizontal plane and temporal distance in the vertical plane, with the past below the cognizer and the future above them. A cognizer's *cognitive reach* — their capacity to take account of distant states of affairs in the generation of its behaviour — can be depicted as the accessible subset of the time/space cone.

Classical cognitive science holds that TSS is possible because behaviour can be influenced by internal surrogates for distant things — viz. mental representations. On this reading cognizers do not use their sensory systems to access information about distant entities but instead rely upon a separate world modelling system that creates internal, 'virtual' environments. These virtual environments provide sensory system-independent input to the behaviour generation process. By contrast, ECS is much more at home with the idea that stretching is enabled by some kind of extended detection system (of which more below).

It is possible that the classical and ECS projects may collapse into one — that what classicists take to be an internal world-modelling system will actually turn out to be realised by some redeployment of detection/action capacities. Be this as it may, I think it is worth persevering with developing the 'stretched detection approach,' because it makes the reasonable evolutionary demand that we explain how animals exhibit knowing about distant entities using brains and bodies that are primarily designed to cope with their local environments. That is, ECS encourages us to pay close attention to the possibility that apparently representation-hungry behaviour (Clark & Toribio, 1994) can be produced through the savvy use of non-neural resources by a situated (that is, non-representational, perception/action-oriented) creature (Anderson, 2007). Although it is beyond the scope of this chapter to attempt a thorough examination of the ways in which a detection-based explanation might make sense of strategies different animals use to access different regions of time/space, a brief first pass will serve to illustrate the nature of the task and the theoretical challenges that it poses.

THREE TIME/SPACE STRETCHING STRATEGIES

Given a view of a cognizer as an entity that strives to access information from the various regions of time/space so that it is well enough informed to behave adaptively, our next task is to try and understand the different ways in which animals might actually do this. What follows is highly speculative and is provided mainly to illustrate the nature of the explanatory task and to make explicit some of the challenges it raises. I thus make no claims to thoroughness here. It would be unwise to both ignore the often surprising TSS capacities of even simple organisms, or ignore the fascinating variations in these capacities among even closely related creatures.

To a first approximation I think it makes sense to distinguish three broad varieties of TSS strategy: first, a largely perceptual strategy utilised by reactive, stimulus-bound

creatures, second a flexible tracking strategy utilised by more complex animals, and third a (probably uniquely human) deliberative strategy. I will consider each in turn.

Reactive, Stimulus-Bound Creatures

If the revolution in situated robotics since the late 1980s has taught us anything it is that quite complex looking adaptive behaviour can be generated by creatures that utilise relatively simple sensory information from their local environments to control the generation of behaviour (Brooks, 1999; Hill, 2006; Pfeifer & Scheier, 1999). Although situated robots are simpler than real insects their behaviour often appears efficient and well-informed. Such creatures inhabit that relatively small region of cognitive time/space that can be accessed by sensory-motor systems with a rudimentary ability to modify the way they work based on a history of interactions with their environment. If one was to mark out the subset of time/space accessed by a situated robot it would not extend much beyond the reach of its sensors.

Importantly though, such creatures *can* access parts of cognitive time/space that lie outside of their immediate sensory world by *tracking* 'out-of-range' targets. Bacteria such as *e. coli* do this by 'swimming' down chemical gradients indicative of food (van Dujin, Keijzer, & Franken, 2006), and female mosquitoes do a similar thing when tracking blood using CO₂ gradients (Christensen & Hooker, 2000). This method of 'knowing' about the presence of out-of-range targets depends crucially upon the existence of environmental regularities. For instance, if high concentrations of CO₂ were not reliably associated with the presence of blood hosts the cognitive apparatus of female mosquitos would not be powerful enough for it to locate blood. Thus, roughly speaking, such environmental regularities form part of a distributed detection system. One might say that the knowledge that blood occurs near CO₂ producing systems is partly in the environment and partly in the 'head' of the mosquito.

Long Distance Pattern Learners

Many animals exhibit a more sophisticated variety of TSS than that of reactive creatures and consequently they inhabit a much larger region of time/space. Researchers are continually revealing the remarkable TSS abilities of different animal species. Some bird species, such as Clark's nutcrackers and scrub jays, exhibit phenomenal abilities to remember where they have hidden caches of food (Suddendorf & Busby, 2003). Many primates possess sophisticated knowledge of the personal qualities of kin and other conspecifics (Clark & Toribio, 1994). Predators, such as cheetahs, learn about the subtle signs associated with the presence of prey and qualities of that prey. All of these creatures know a great deal about distant environments such as where they buried their seeds, what will likely happen should they attempt to steal food from a higher ranked monkey, and what might lie in wait at the waterhole. These animals differ from reactive creatures in two important ways: 1) they are sensitive to the distant consequences of a multitude of cues in the local environment, and 2) they can learn new associations between local environment cues and distant environment opportunities and threats. Clearly the behavioural flexibility of these animals is far beyond anything exhibited by reactive creatures and clearly such flexibility demands complex cognitive mechanisms.

The question is, can these sophisticated abilities be explained by a detection-action framework?

It is tempting to argue that increased ability to access distant information demands a representational explanation — if the environment cannot unambiguously inform behaviour generation then the mind (mental representations) must do it instead. Writers, such as Christensen and Hooker (2000) and R. Wilson (2004) argue that sophisticated cognitive abilities imply a greater internal control over the deployment and modification of detection-action systems than is found in simpler reactive creatures. While I am sure that long distance pattern learners *do* have greater internal control (in the sense that they are not completely at the mercy of the environment as many reactive creatures are), I think it would be a mistake to think that this increased complexity and self-sufficiency somehow implies that these animals are any less embedded in, and reliant on, their environments for their cognitive achievements. On the contrary the ‘worlds’ of these creatures (their *Umwelts*) are larger, more complex, and more dynamic than those of simple creatures and thus, in an important sense, long distance pattern learners are *more* reliant on their environments for producing their fine-grained, varied, and robust behaviours. Whereas an *e. coli* bacterium lives in a simple world of nutritious medium and non-nutritious medium, a cheetah must distinguish thousands of different cues that signal resources and dangers that impact on a complex set of needs. Long distance pattern learners need bigger and more sophisticated internal control systems to cope with (and enact) their larger and more complex *Umwelts*. So instead of thinking that complex abilities for gathering information about distant events and objects requires (decoupled, offline) representations of those objects (because they are perceptually out-of-range) one can argue that long distance pattern learners are very good at learning how subtle signs in the local environment predict distant targets. We should avoid the temptation of thinking that more complex nervous systems *necessarily* imply more ‘representational’ or ‘off-line’ minds and instead understand them as providing a substrate for complex, sensitive, and subtle capacities for detection and action.

In many ways the TSS capacities of long distance pattern learners are broadly similar to those of reactive agents. They too access the distant parts of cognitive time/space by tracking gradients — that is by, in a sense, seeing signals of the past, the present, and the out-of-range in the present (the here and now). These signals are unlikely to arouse *recollections* of the distant associated targets (i.e., true episodic memory), rather they simply engender a generalised sense of needing to modulate behaviour in anticipation of the currently out-of-range target (although see the debate between Suddendorf & Busby, 2003 and Clayton et al., 2003). Long distance pattern learners differ from reactive creatures in that they actively learn how to ‘see’ the world in new and more useful ways (including learning to ‘stretch farther and more reliably’ by finding new ways to track out-of-range targets) in order to better produce adaptive behaviour. Explaining exactly how such long distance pattern learning may be possible is beyond the scope of this chapter but Keijzer (2001) provides a sketch of how such a system might work.

Human Imaginative Access to Distant Regions of Time/Space

One might think that the enormous potential flexibility of the long distance pattern learner strategy is enough to explain the TSS abilities of humans, but these capacities pose a number of puzzles for cognitive scientists that I believe require an additional 'stretching strategy'. As clever as many non-human animals are they do not seem to possess the high levels of flexibility of interpretation and the ability to 'think' in a largely context-independent and offline fashion that (adult?) humans do. Put crudely, we humans can pretty much sit down anywhere and think about anything we like – we do not seem to rely on 'environmental stimuli' to drive all, or even much, of our behaviour (M. Wilson, 2002). In addition to the methods available to reactive creatures and long distance pattern learners, humans possess complex cognitive abilities that enable us to 'think about' distant regions of space and time including hypothetical or imaginary versions of our world. Indeed I think that it probably this strong type of imagination that underpins our unique stretching abilities (see Thomas, 1999 for a related idea). Imaginative processes enable us to flexibly move forward and back in time, and to parts of the present that we cannot access in a seemingly unconstrained fashion. In doing this we seem to open unlimited access to 'information' that can be used to modulate our behaviour (see Suddendorf, 1999 for a related idea). For instance, in prediction humans do not just anticipate how things will turn out in the future, they actively imagine them – in a sense mentally travelling in time (Suddendorf & Corballis, 2007). Whereas many non-human animals might exhibit cautiousness about what is likely to happen when they detect signs of danger, it is unlikely they entertain full blown 'mental simulations' of what they will find, and even more unlikely that they can 'run them' at their leisure. As we all know this kind of thinking is common in humans. Similar considerations arise when thinking about the ways in which humans and non-humans can access aspects of the present that are out of perceptual range, and for thinking about the past. Indeed knowing about the past — in the sense of possessing episodic memory — is arguably a uniquely human ability (Suddendorf & Corballis, 2007). Nonhumans' contact with the past may only be evident in the ways in which current behaviour is shaped by past experience.

These imaginative capacities have been discussed at length by numerous theorists and they strongly encourage the metaphor that the human mind must incorporate an internal theatre of computation that manipulates free floating representations of states of affairs (e.g., Amati & Shallice, 2007). It is widely assumed that human cognitive and behavioural flexibility cannot be a product of an 'environmentally driven' system and thus must be generated by an independent in-the-head system that can create its own virtual stimuli on demand. Cognitive scientists have found it tempting to suggest that, although a detection-action architecture can account for our ability to respond to many immediate threats and opportunities in the local environment, our more reflective and deliberative abilities require a more traditional kind of 'world modelling' architecture. Thus, there are significant challenges for anyone interested in constructing a convincing detection-action explanation of the human capacity for accessing the entirety of cognitive time/space. What might a detection-based theory of human imaginative time-space stretching look like? The vision I have goes something like this: humans have a capacity for engaging with distant environments via a kind of perceptual bootstrapping in which we examine an environment and see that it tells us something useful *including*

how to create 'new epistemic objects' (NEOs) — words, symbols, artifacts, whatever — that can be used to tell us something about the state of distant environments. These NEOs can then serve as stimuli to be examined *in the context of the local environment*. That is, we change the local environment so our detection capacities can tell us more about distant environments.

Here is a simple, trivial example: You are waiting to go into a job interview but do not know what it will be like. Another applicant who has already been interviewed returns to the waiting room (the local environment). You ask them what the interview was like (an epistemic action). Their verbal response (a NEO) changes the local environment and your perception of it tells you something about the nature of your future interview (the target distant environment) providing, of course, that the applicant responds, responds truthfully, and provides useful information.

In other situations we run through this entire process privately (although, à la Vygotsky, we may have originally learnt how to do this in a public, social context). In these cases, the NEOs are words, images, or other symbols, that we might say out loud to ourselves, or write down, or sketch out. That is, these NEOs are *material symbols* (Clark, 2006, Roepstorff, 2008) and they provide fuel for detection-action systems as well as being the *products* of earlier detection-action cycles. Adult humans often do not need to actually produce the material symbols to access distant environments — subvocalisation or 'shallowly encoded mental images' (Clark, 2006) will do (although there are substantial advantages in producing material symbols over 'doing it all in your head', especially with regard to complex, future-oriented activities [see e.g., Clark, 1997]). That is, our bodily detection-action systems can be activated so that they are in a state necessary for producing NEOs but without the requisite efferent activity that leads to the modification of the environment. A good deal of research supports the idea that imagining words, images, or actions involves neural processes and structures that also give rise to overt action and symbol production (see Roepstorff, 2008; van Dijk et al., 2008). Thus, in this view the local worldly environment (outside of the body) does not necessarily change but the detection-action system changes as if it had. Moreover, none of this would be possible if we did not regularly produce material symbolic NEOs. Imaginative access to distant environments is not simply an ability to directly perceive those environments because it requires the active epistemic manipulation of the local environment. However, neither is it the subpersonal model-building, in-the-head activity described in many a cognitivist theory. Rather, stretched detecting is enabled by the canny manipulation of the local environment so that some knowledge of distant environments can be accessed.

Clearly for such a system to work it has to be principled. We must possess a 'grammar' of 'NEO building' that ensures that the NEOs we construct tell us something useful about distant environments. Sometimes NEO building will not work. Often it will *partly* work — perhaps well enough to lead to the production of slightly more adaptive action than would have been the case without it. Sometimes things might go badly wrong and our actions will be much less adaptive than would have been the case using the strategies found in long-distance pattern learning. Any developed theory of stretched detection will need to include a sophisticated account of how we learn to constrain our NEO building. This is a particularly challenging task because human imagination is, by

its very nature, hugely, perhaps infinitely, flexible. That is, we have the capacity to imagine many kinds of distant environments *including those that are not intended to be accurate renditions of any actual distant environment*. How is it that we can do the latter while at the same time constraining some of our imaginative thinking to helping us understand ‘real’ distant environments? How is it that we come to trust some of our imaginings more than others? What function, if any, do our ‘fictional imaginings’ have? Are they simply a (mostly) pleasant byproduct of the mechanisms we possess for understanding ‘real’ distant environments or do they in fact aid our understanding of distant environments (possible futures, new ways of doing things) in indirect ways? All of these questions, and more, face the stretched detection theorist.

WHAT SHOULD BE EASY AND WHAT SHOULD BE DIFFICULT WHEN EXPLAINING COGNITION?

Explaining the human capacity for ‘imagination-based’ time/space stretching using a detection-action framework is going to be a complex and difficult task. Some might think that the major challenges we face trying to develop a detection-action explanation of complex ‘stretching’ strongly suggests that the detection framework cannot be scaled up to explain advanced human cognition. However, it seems to me that this is exactly what we should expect and exactly what we do not get from classical cognitive science. Uniquely human cognition is profoundly unusual and puzzling when set against the capacities of other animals. So one should expect that there is something unusual about the ways in which humans deploy their neural and extra-neural resources that makes this so. Taking the classical cognitive science route and simply appealing to a representational mind, which is also used to explain all other cognitive activity, is not enough. At best the traditional view that “it’s representations all the way down” obscures the complexity of the explanatory task we face; at worst it might misleadingly portray the representational framework as having already successfully explained advanced human cognition. Thus, the fact that the detection-action approach demands extensive further theorizing may better represent our current understanding of the ways we stretch time and space.

REFERENCES

- Amati, D., & Shallice, T. (2007). On the emergence of modern humans. *Cognition*, 103, 358-385.
- Anderson, M. L. (2003). Embodied cognition: A field guide. *Artificial Intelligence*, 149, 91-130.
- 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.
- Barsalou, L. W., Breazeal, C., & Smith, L. B. (2007). Cognition as coordinated non-cognition. *Cognitive Processing*, 8, 79-91.
- Brooks, R. A. (1999). *Cambrian intelligence: The early history of the new AI*. Cambridge, MA: MIT Press.
- Christensen, W. D., & Hooker, C. A. (2000). An interactivist-constructivist approach to intelligence: Self-directed anticipative learning. *Philosophical Psychology*, 13, 5-45.
- Clark, A. (1997). *Being there: Putting brain, body, and world together again*. Cambridge, MA: MIT Press.
- Clark, A. (2006). Material symbols. *Philosophical Psychology*, 19, 291-307.

- Clark, A., & Chalmers, D. J. (1998). The extended mind. *Analysis*, 58, 7-19.
- Clark, A., & Toribio, J. (1994). Doing without representing? *Synthese*, 101, 401-431.
- Clayton, N. S., Bussey, T. J., Emery, N. J., & Dickinson, A. (2003). Prometheus to Proust: the case for behavioural criteria for 'mental time travel.' *Trends in Cognitive Sciences*, 7, 436-437.
- Hesslow, G. (2002). Conscious thought as simulation of behaviour and perception. *Trends in Cognitive Sciences*, 6, 242-247.
- Hill, S. (2006). Are humans mobot minded? Some implications of embodied cognitive science for cognitive psychology. In C. M. Fletcher-Flinn & G. M. Haberman (Eds.), *Cognition, language, and development: Perspectives from New Zealand* (pp. 257-269). Brisbane, Australia: Australian Academic Press.
- Keijzer, F. (2001). *Representation and behavior*. Cambridge, MA: MIT Press.
- Kosslyn, S. M. (2006). On the evolution of human motivation: The role of social prosthetic systems. In S. M. Platek, T. K. Shackelford & J. P. Keenan (Eds.), *Evolutionary cognitive neuroscience* (pp. 541-554). Cambridge, MA: MIT Press.
- Pfeifer, R., & Scheier, C. (1999). *Understanding intelligence*. Cambridge, MA: MIT Press.
- Roepstorff, A. (2008). Things to think with: words and objects as material symbols. *Philosophical Transactions of the Royal Society B*, 363, 2049-2054.
- Suddendorf, T. (1999). The rise of the metamind. In M. C. Corballis & S. E. G. Lea (Eds.), *The descent of the mind: psychological perspectives on hominid evolution*. (pp. 218-260). Oxford: Oxford University Press.
- Suddendorf, T., & Busby, J. (2003). Mental time travel in animals? *Trends in Cognitive Sciences*, 7, 391-396.
- Suddendorf, T., & Corballis, M. C. (2007). The evolution of foresight: What is mental time travel, and is it unique to humans? *Behavioral & Brain Sciences*, 30, 299-351.
- Sutton, J. (2006). Distributed cognition: Domains and dimensions. *Pragmatics and Cognition*, 14, 235-247.
- Thomas, N. J. T. (1999). Are theories of imagery theories of imagination? An active perception approach to conscious mental content. *Cognitive Science*, 23, 207-245.
- van Dijk, J., Kerkhofs, R., van Rooij, I., & Haselager, P. (2008). Can there be such a thing as embodied embedded cognitive neuroscience? *Theory & Psychology*, 18, 297-316.
- van Dujin, M., Keijzer, F., & Franken, D. (2006). Principles of minimal cognition: Casting cognition as sensorimotor coordination. *Adaptive Behavior*, 14, 157-170.
- Wheeler, M. (2005). *Reconstructing the cognitive world: The next step*. Cambridge, MA: MIT Press.
- Wilson, M. (2001). The case for sensorimotor coding in working memory. *Psychonomic Bulletin & Review*, 8, 44-57.
- Wilson, M. (2002). Six views of embodied cognition. *Psychonomic Bulletin & Review*, 9, 625-636.
- Wilson, R. A. (2004). *Boundaries of the mind: The individual in the fragile sciences – Cognition*. Cambridge, UK: Cambridge University Press.
- Wilson, R. A. (2007, June 18th). *The extended mind*. Keynote address presented at the International Society for Theoretical Psychology Conference, Toronto, Canada.