

Differentiating Animality from Agency Towards a Foundation for Cognition

Fred Keijzer (f.a.keijzer@rug.nl)

Department of Theoretical Philosophy,

Oude Boteringestraat 52, 9712 GL Groningen, The Netherlands

Abstract

The notion of cognition has been difficult to pin down. Embodied and situated approaches to cognition now suggests that agency, construed in terms of perception-action coupling, might provide a clear foundation for cognition. Yet, this attempt has problems of its own. First, a demarcation problem: Which systems are agents and why? Second, a graduality problem: Agency is a way of describing that is either used or not, and it is difficult to envision an incremental route toward agency. To overcome these problems, I differentiate between agency and a new notion, animality. Animality can be described as the sensorimotor organization by means of which animals modify environmental conditions. By developing this notion of animality I claim that it becomes possible to get a grip on the foundational problems of cognition and agency.

Keywords: Cognition; agency; animality; perception-action; intentional stance; philosophy; biology; cognitive science.

Introduction

“There is actually little consensus as to what makes something a cognitive process. Often, the class of processes that we regard as cognitive is defined by ostension.” (Rowlands, 2003, p.157). Mark Rowlands nicely summarizes here both the current lack of clear ideas about what makes something a cognitive system, as well as the lack of a feeling of urgency when this is brought to the fore. We simply point to processes like perceiving, remembering, thinking, reasoning and language and take these to make up the cognitive domain. In addition, the standard cognitive science view is that almost any system can be considered cognitive as long as we think it is usefully described in these cognitive terms. Beyond that, it does not seem to matter very much whether they are physically computers, animals, or humans. Nevertheless, it remains unclear why or when these entities can be deemed cognitive. To illustrate the point, 56 years after Turing’s first try at an answer, we still have no clear criteria to decide whether an AI program has to be considered a *model* or an *instance* of a thinking system. The question: “What makes something a cognitive system?” is far from solved.

After a long time of neglect, the question currently gains a new urgency as embodied and situated cognition now uses the notion of cognition in different ways that require clarification. Also, the standard way of demarcating the cognitive domain is unsatisfactory, as I will argue below.

In this paper, I address the question posed above by introducing the concept of *animality*; the sensorimotor organization by means of which animals modify environmental conditions. The aim of introducing animality is to demarcate a natural domain, which is plausibly cast as the material foundation of cognitive systems. First, I will discuss why the current delineation of cognition is problematical; second, how embodied and situated cognition tries to base cognition in agency, and why this project is hampered by similar problems; third, in order to overcome these problems, I differentiate between agency and animality and develop the latter in a preliminary way. Finally, I will conclude that animality is a plausible candidate as a material foundation for cognition and agency.

Problems with Ascribed Cognition

Within cognitive science there is a strong tendency to ascribe cognition when we can interpret a system in terms like perceiving, remembering etc. Dennett (1981) introduced the notion of an *intentional stance* to describe this way of having intentional—mental or cognitive—systems. Taking an intentional stance amounts to treating a system as a rational agent and figuring out what its beliefs and desires are likely to be. In this way, we can often understand and predict what such a system will do without needing to know about its detailed physical makeup. The intentional stance provides a way to combine our mental or intentional vocabulary with a mechanistic understanding of cognitive phenomena. One can use the mental vocabulary to predict and explain particular, cognitive, systems, while the stress on it being a *stance* makes it perfectly clear that this is merely a different description of a physical system. There is no risk of an unaccounted for ghost in the machine.

Cognition thus conceived is a useful and pragmatic way of demarcating the cognitive domain. Nevertheless, it comes at great theoretical costs. Most notably, I will argue, is that it obstructs a clear linkage between the cognitive domain and particular kinds of material systems.

A Double Bind

Suppose that systems are to be considered cognitive or not on the basis of applying an intentional stance: If such a stance is applied successfully, then the system counts as a cognitive one. However, as we are free to apply the intentional stance to whatever we want—be it a falling stone, thermostat, animal or human being—and given that the notion of successful appliance is open to many different interpretations, this is a very unconstrained way of

demarcating a cognitive domain. In particular, it amounts to a way of having cognition which remains independent of any particular material organization of the so described systems. Thus the view arises that, even though an AI program runs on a computer and not 'on a brain,' it still may be considered cognitive.

But there is also an opposing intuition: There ought to be something about the systems themselves that makes them cognitive or not: Humans and falling stones are just too different. The issue also comes to the fore in research on animal cognition, where it is a major research effort to establish whether particular animals can be deemed cognitive or not. Thus, there is clearly more to be said about delineating cognitive systems than applying an intentional stance. According to this intuition, we are not free to postulate cognition wherever we want.

The way to go from here is investigate the physical means that might account for the differences between genuine cognitive systems and systems more generally. However, an intentional stance deemphasizes the means by which an agent achieves its goals. Dennett once introduced the revealing phrase 'wise wiring' (1987), which illustrates the point. The phrase 'wiring' refers to an arbitrary set of connections that are in themselves neither systematic nor very important as long as it produces the required 'wise' result. This choice of words implies that the physical system involved is not special but an *ordinary* system, which merely happens to produce particular results. Thus, if one would stress the importance of the wiring, then this would count as a dismissal of the need for cognition, rather than its articulation. In addition, tying cognition to particular kinds of systems will always exclude systems that are currently taken to belong to the cognitive domain as derived from the intentional stance. Of course, most of those exclusions would be according to the intent of such an endeavor, but border disputes would nevertheless arise, which might seem to discredit the very project.

These opposing tendencies lead to a double bind when it comes to answering the question of what cognition is. The intentional stance provides an intuitively plausible cognitive domain, but remains too unspecific to be the whole story. At the same time, the force of the intentional stance criterion makes it almost impossible to delineate cognitive systems in a more specific way, because it is not the 'wiring' that counts, but its being wise.

A Conceptual Dichotomy

In summary, an intentional stance sets up a *conceptual* dichotomy between intentional and mechanical systems, but does so without a corresponding dichotomy between different kinds of material systems, and even prohibiting any material distinction to be the relevant one. At the same time, there are good reasons to suppose that there must be more specific physical, organizational or dynamical aspects to cognitive systems that ought to set them apart as a particular kind of material system. Thus, taking cognition as something that can be simply ascribed by taking an intentional stance may suffice for cognitive science in the

short run, but it also obstructs raising the question what cognition could be on a more fundamental material level.

Founding Cognition in Agency?

How to proceed? In the last fifteen years or so, embodied and situated interpretations of cognition have been critical of interpreting cognition in terms of internal reasoning processes. Instead, they made a strong case for putting cognition squarely in the context of perception-action relations (Brooks, 1999; Clark, 1997; Hurley, 1998; Keijzer, 2001; Pfeifer & Scheier, 2001; Van Gelder, 1998). Embodied and situated cognition does not claim that all cognitive processes consist only of perception-action relations, but it does tend towards the view that even presumably fully internal cognitive processes as remembering and reasoning are ultimately based in perception-action systems. Without going into these details here, in my view, embodied and situated approaches have promising implications for developing more specific answers to the question what cognition might be.

Firstly, embodied and situated approaches place cognition in the context of agents who perceive and act in an environment. This forms a firm step to a more concrete interpretation of cognition. Secondly, a perception-action interpretation is more congenial to a biological and evolutionary perspective on cognition. When primarily interpreted as inner thought, cognition remains almost specifically human and its link with biology and evolution is not self-evident. In contrast, perception and action are spread widely across the biological domain and have a clear evolutionary relevance. Such a link with biology is good because biology provides much stronger and more detailed constraints on cognition than what can be derived from the mental vocabulary (Lyon, 2006).

Despite these positive aspects, there are important problems associated with grounding cognition in agency. I will discuss two of them.

A: *The demarcation problem*: When we take agency as the deciding factor for considering systems cognitive or not, what are the criteria to distinguish agents from non-agents? One should get a *déjà vu* here: Do falling stones, computers, animals and software agents all provide examples of perception-action systems? Clearly, when the notion of cognition is difficult to pin down, the same problem applies to agents (Wooldridge & Jennings, 1995).

B: *The graduality problem*: How can this notion of agency be cast in an incremental way? Being considered an agent seems to be an all or nothing affair. The famous analytical philosopher Donald Davidson notes: "We have many vocabularies for describing nature when we regard it as mindless, and we have a mentalistic vocabulary for describing thought and intentional action: what we lack is a way of describing what is in between. This is particularly evident when we speak of the 'intentions' and 'desires' of simple animals; we have no better way to explain what they do" (1999, p.11). Of course, we can easily think up a gradient of very stupid to very smart agents, but then we do

not have a gradient from the mindless to the mindful. Even a stupid agent is a full agent and we still have the problem of how to reach the point of such minimal agency in a gradual and non-arbitrary way.

The analysis of these problems is straightforward: Agency is also a matter of ascription, and itself part and parcel of the intentional stance. Agency, then, does not provide a material foundation for cognition. However, it does form a signpost in the right direction.

Differentiating Animality from Agency

Agency promised to contribute two positive features to a possible material foundation for the notion of cognition—perception-action relations and a biological context—but could not deliver the goods. To proceed, a different provider is necessary. The crux to progress, is to *disregard the intentional stance criterion* altogether and to turn directly to those physically constituted systems that definitely embody the combined characteristics of perception-action relations as well as a biological context.

Complying with the second characteristic is relatively easy because one can simply refer to living organizations, which have a clear scientific status. One can even leave it at that and take life itself as designating the cognitive domain (Maturana & Varela, 1980; Stewart, 1996). However, it seems preferable to cast cognition in terms that also take the perception-action aspect as a precondition (Van Duijn, Keijzer & Franken, in press). As I have just discarded the intentional idiom as a backdrop for concepts like perception and action, it is also essential to provide a different foundation for these terms.

Both desiderata can be had by turning to animals, or rather animalia.¹ These are concrete living systems where the notions of perception and action readily apply. To mark the difference in background, I will use the phrase *sensorimotor relations* henceforward. Animalia constitute a set of systems that is markedly smaller than life itself, includes the human case, and involves systems at widely varying levels of sensorimotor complexity. Suppose now that we take animalia as the paradigmatic cognitive systems, then, the key question becomes: What is it about animalia that makes them cognitive? The answer cannot be formulated in terms of agency, because this is precisely the notion that we seek to find a foundation for.

In the following, I propose a principled distinction between the notion of agency and what I call *animality*. Animality refers to the structural dynamical sensorimotor organization by which animalia modify environmental conditions through movement. These result in dynamical relations that embody particular fleeting dynamical structures which subserve the metabolic and reproductive requirements of the living organization to which they belong.

¹ Animalia because being free-moving creatures is what counts, sessile animals being a borderline case, while free-moving bacteria as well as protozoa should fall in.

While the notion of agency, and cognition, can be applied very widely, animality is restricted to animalia, and refers to the specific organizational setup responsible for generating the agentive characteristics exhibited by these systems. Animality does not depend on an intentional stance but aims to articulate a set of mechanisms (Bechtel & Richardson, 1993) that together give rise to behavioral-cognitive phenomena across the animal kingdom.

Animality

The notion of animality provides a relatively unencrusted term that allows one to focus on those aspects of animalia's sensorimotor organization that tend to be obscured by an agentive terminology. In the following, I will try to clarify the notion of animality by discussing its derivation from the work of Hans Jonas, by providing an example that draws out the contrast with agency, and, finally, a first try at a more detailed analysis.

“To Move and to Feel”

Hans Jonas used the phrase animality in a book that aimed to link human existence to biology (1966).² In one essay, “To move and to feel,” Jonas describes the switch from plant-life to animal-life. “Three characteristics distinguish animal from plant-life: motility, perception, emotion.” (p.99) Emotion comes in because motility induces a way of life that breaks the immediate and reliable organism-environment relations of plant-life. To gain access to metabolically necessary nourishment, animal-life builds on a fickle, spatiotemporally drawn out process involving multiple steps. “The very span between start and attainment which such a series represents must be bridged by continuous emotional intent.” (p.101) Jonas used the phrase animality only in passing to refer to this characteristic setup for animal-life, but the notion targets precisely the motile and sensory setup which allows animalia to thrive.

The obvious aspect of animal motility is that it enables the creature to move itself and to manipulate its environment. What I want to stress here is the specific organizational makeup behind motility that is very different from robot-effectors (Keijzer, 1998; Sharkey & Ziemke, 2001). Particularly for multicellular creatures, the capacity to move is not a primitive, but in itself a significant achievement which requires the generation of patterns across the body—for example undulations, locomotory waves or leg movements—and which is totally wrapped up with the specific characteristics of the body, e.g. whether the animal has a soft or hard skeleton, number of appendages and so on. To be motile for an animal requires a complex pattern generation process, which from an agentive perspective is easily left out of consideration. The role of the nervous system, if present, also becomes more easily cast in terms of pattern generation across the effector surfaces rather than executing relatively abstract tasks.

² Jonas's work was brought to my attention by a review of Di Paolo (2005).

Sensing in animalia is likewise a matter of being sensitive to patterns of change on sensory surfaces, whether these are chemical, tactile, electromagnetic or other. Again the particulars of the sensory surfaces are crucial to understand what is going on, and to understand how the animal operates. Questions that immediately arise in this perspective is how sensory and effector patterns relate to one another, and how the nervous system, if present, fits in.

To wrap up, animality refers to the detailed structures and their role in the dynamical sensorimotor processes that are at work in animalia, and which are intrinsically related to metabolic functioning. I will now use the jellyfish *Aglantha digitale* to illustrate the contrast with agency.

A Hydromedusan Example³

The hydromedusan jellyfish *Aglantha digitale* consists of a transparent bell, which has an opening at the base. Movement of the creature is achieved by the patterned contraction of muscles set across the bell, making the bell itself contract, pushing water outside through the opening and so providing a kind of jet propulsion. From the margin at the base of the bell, many fine tentacles extend outwards. When small planktonic creatures touch these, they are killed by independently acting stinging cells (nematocysts) and carried to the margin by tentacle flexions (Mackie et al, 2003). Subsequently the manubrium, say the mouth, bends toward the prey and engulfs it (*ibid.*).

Taken as an agent, *Aglantha* does not amount to very much. Disregarding its feeding behavior and other intricacies, *Aglantha* does two things, swimming slowly to feed and, when touched, escaping by a fast swim. Dennett in his inventory of different kinds of mind would designate it as a *Darwinian creature* (1996, p.110), situated at the ground floor of his hierarchy as a creature hardwired by evolution about which nothing much needs to be said from a cognitive perspective. Similarly, Sterelny in *The evolution of agency* targets in this context “the evolution of belief-like states, and the evolutionary transition from organisms that detect and respond to their environment in very simple ways to more complex representation by an organism of its environment” (2001, p.21). Again, *Aglantha*, if taken as an agent, would be no more than a starting point that can be described as “very simple.”

Animality provides a different view of the same creature. Foremost, it stresses that *Aglantha* is a living animal with a particular metabolic and cellular organization. This is important because, when compared to the cellular level, *Aglantha* is a huge organization which involves different and complex new forms of coordination compared to those on the cellular level. The animality present in *Aglantha* is thus not its simple agentive functionality—deciding to swim either fast or slow—but, rather, the kind of problems that must be overcome to produce such large-scale functionality, given the initially microscopic cellular building blocks. As

an analogy, one may consider the cognitive task of building an arch from three bricks as fairly trivial, but this changes radically when the bricks transform into huge megaliths weighing tons. The abstracted description of this task does not give sufficient guidance concerning its actual difficulty, as the latter is relative to the means available. An animalistic perspective on *Aglantha* then targets the processes that bring about behavioral functionality, instead of an agentive perspective which merely assumes that there is a set of processes that does the trick. The notion of animality draws attention to neural and sensorimotor means involved, and these are highly complex for all animalia. To press this point, given that biologists readily acknowledge the complexity of cellular signaling in biology, cognitive scientists should not hesitate to do the same when it comes to the cellular and neural signaling processes that occur in a nervous system as present in *Aglantha*.

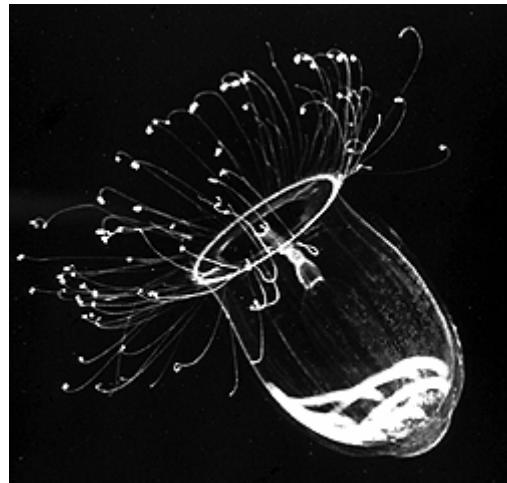


Figure 1: *Aglantha digitale* showing the bell and the tentacles attached to the margin. Right in the middle of the margin, the manubrium or mouth is visible.
(Photograph by Claudia E. Mills).

A Conceptual Elaboration

The *Aglantha* example may seem to convey the message that animality amounts to a neuroethological view, targeting the details of the behavioral and neural mechanisms in comparatively simple animals. What does the notion of animality add to this ongoing enterprise? First note that the concept of animality is brought forward to help answering the question what cognition is, and not as a different empirical approach to these behavioral and neural mechanisms. Animality should build on existing work as done within neuroethology and other relevant fields. Having said this, it might nevertheless be that the ongoing study of neural and behavioral mechanisms could benefit from work and concepts from embodied and situated cognition. This might be helpful to arrive at a more cohering global picture of the organizational principles involved in animality. In this section, I will discuss three concepts that could make a difference to ongoing empirical work.

³ This section is based on information drawn from Brusca and Brusca (1990), Mackie, Marx and Meech (2003), Meech (1989), and Singla (1978).

O'Regan and Noë (2001), relying on the classic work of Gibson and others, introduced the notion of *sensorimotor contingencies*. O'Regan and Noë describe these as “the structure of the rules governing the sensory changes produced by various motor actions” (2001, p.941). For example, “when the eyes rotate, the sensory stimulation on the retina shifts and distorts in a very particular way, determined by the size of the eye movement, the spherical shape of the retina, and the nature of the ocular optics” (ibid.). The relations between motor output and sensory input obey specific lawful regularities and provide an important constraint and fundamental principle. The concept of sensorimotor contingencies can be used as a way to systematically investigate the relations between the specific motility and sensing capacities of particular animals and to search for and categorize regularities. As neural systems have evolved to coordinate such sensorimotor contingencies, they could provide an important handhold for their investigation.

Another important concept can be termed *layering*, as e.g. exemplified in Brooks' subsumption architecture for robots (Brooks, 1999) and Hurley's (1998) notion of horizontal modularity. Layering implies that a control structure must first and foremost allow for ongoing sensorimotor relations. Everything else, including human thought, is derivative. Getting a minimal layer of sensorimotor coordination is the prime directive. From there on, improvements are possible by changing or adding to the sensors, effectors or neural systems of the existing system. The key issue with layering is that additional layers are *additional* and not separable from an underlying sensorimotor basis. This iterative buildup also gives a reason to think of animality as a coherent whole, where complex instances of animality form an organic organization, rather than an arbitrary collection of sensorimotor relations that can be freely taken apart.

A third important concept, spatiotemporal *pattern generation*, has already been introduced as a key feature of generating motility. Both behavioral and neural pattern generation is important in neuroethological explanations. A wide variety of behaviors are generated by neural rhythmic pattern-generation circuits. “These include ongoing and stereotyped movements such as breathing, chewing, walking, running, flying, and swimming” (Marder & Calabrese, 1996, p.688). In addition, the sensory shaping of motor patterns is essential as well. “The dynamic interplay between central and sensory mechanisms in the generation of adaptive movements is seen in all preparations” (ibid.). Pattern generation provides a close conceptual link, going both ways, between neural and sensorimotor phenomena. As such, pattern generation and its role in animality could be a way to unravel neural functioning and its relation to sensorimotor phenomena at a more fundamental level and in greater detail than has so far been possible.

In all three cases, it seems that the field of embodied and situated cognition, and neuroethology could be mutually enriching to a greater extent than has so far been the case, leading to a better understanding of animality.

And Human Cognition?

So far nothing has been said about human cognition. This was deliberate as the focus was on animality and the foundation of cognition. In this picture, human cognition is just one case among many, rather than the center of the cognitive domain. Of course, it is legitimate to be primarily interested in human cognition, and it goes without saying that human cognition is hugely different from what happens in *Aglantha*. However, as a general practice in science, for example human genetics, it is unusual to try to tackle the hardest case directly. From this perspective, it goes without saying that one must study more basic cognitive organizations to understand how they work.

Thus, a final strong difference between agency and animality is that agency ultimately provides a human centered perspective—humans being the prime targets for ascribing rationality—while animality puts us in our animal context. We may be very different from even the great apes, but before we can truly understand those differences we should become more sensitive to the huge overlaps between human cognition and that of other animals.

Founding Cognition in Animality?

Summing up: Animality refers to the sensorimotor organization present in animalia, it does not build on agency but provides the kind of organization to which agency can be ‘properly’ ascribed. As yet, animality remains a preliminary notion that can and should be filled in by further research. The question to turn to now is: Can animality provide a suitable foundation for cognition? In the following, I will return to the two problems that agency encountered when cast in this role, starting out with the good news and then seeing whether there is any bad.

First, animality provides a clear solution to the demarcation problem. Animalia are cognitive systems in some form or other, thermostats, robots and computers are not. The latter have a different kind of organization and for this reason cannot be deemed cognitive. This demarcation also seems sufficiently principled. Sensorimotor relations are plausibly the starting point of all animal cognition, while, at the upper level, they even generate ideas that might help explain how the brain gives rise to consciousness (Hurley & Noë, 2003).

Second, animality also provides a solid way to deal with the graduality problem. The concepts of cognition and agency are notions that are specified irrespective of a particular physical organization, concerning which they are taken to apply, or not. Such a background makes a gradient from the physical to the cognitive problematical. Animality, in contrast, is a particular kind of physical setup, which during evolution arose first in a basic bacterial form, and from there on developed into many different forms, some of which are hugely more complex, like the human case. Graduality is part of the notion of animality from the very start.

So far for the good news, what about the bad? Actually, I think there isn't much, even though it may seem like that. One seeming problem might be that the animality criterion

cuts of too many plausible cases of cognition, such as in AI or robots, another one could be the seeming lack of applicability to genuine human cognition and consciousness. As space is extremely limited, I will just hint at the kind of answers that can be given here.

Are AI systems and robots wrongly left out? Let me just use an analogy: Could biology be criticized for leaving out of consideration Artificial Life models as clear cases of life? I think not. The differences are too huge. ALife models are life-like but, at present, not yet living themselves. I would argue that the same holds for the relation between AI and cognition (see also Sharkey and Ziemke, 2001).

Is human cognition insufficiently dealt with? Again, no. The current project is to locate human cognition within the general natural science picture, not to provide an account of human cognition itself. It goes without saying that human cognition goes way beyond the simple cases that received attention here. However, to really understand the human case, it must be considered essential to understand the operation of nervous systems more generally as well as how this operation is linked up with sensorimotor processes, the *raison d'être* for any nervous system. It seems an irresponsible procedure not to pay close attention to simpler case studies, even when the ultimate goal is strictly human cognition.

To conclude, I hold that animality provides a plausible articulation of the material kind of systems to which notions like cognition and agency most readily apply. From here on, we might start to consider what the implications are for these notions and for cognitive science.

Acknowledgments

I thank Daan Franken, Marti Hooijmans, Barteld Kooi, Erik Krabbe, Theo Kuipers, Pamela Lyon, Allard Tamminga, Marc van Duijn, members of the Groningen Theoretical Philosophy research colloquium, and five anonymous reviewers for their comments and/or discussion. This work was supported by the NWO research grant 016-038-301.

References

Bechtel, W. & Richardson, R.C. (1993). *Decomposition and localization as strategies in scientific research*. Princeton, NJ: Princeton University Press.

Brooks, R.A. (1999). *Cambrian intelligence*. Cambridge, MA: MIT Press.

Brusca, R.C. & Brusca, G.J. (1990). *Invertebrates*. Sunderland, MA : Sinauer.

Clark, A. (1997). *Being there: Putting brain, body and world together again*. Cambridge, MA: MIT Press.

Davidson, D. (1999). The emergence of thought. *Erkenntnis*, 51, 7-17.

Dennett, D.C. (1981). *The intentional stance*. Cambridge, MA: MIT Press.

Dennett, D.C. (1987). Eliminate the middleload! *Behavioral and Brain Sciences*, 10, 372-374.

Dennett, D.C. (1996). *Kinds of minds: The origins of consciousness*. London: Phoenix.

Di Paolo, E.A. (2005). "The phenomenon of life" by Hans Jonas. *Journal of the British Society for Phenomenology* 36(3), 340 - 342.

Hurley, S.L. (1998). *Consciousness in action*. Cambridge, MA: Harvard University Press.

Hurley, S.L. & Noë, A. (2003). Neural plasticity and consciousness. *Biology & Philosophy*, 18, 131-168.

Jonas, H. (1966). *The phenomenon of life*. New York: Delta.

Keijzer, F.A. (1998). Some armchair worries about wheeled behavior. In R. Pfeifer, B. Blumberg, J-A. Meyer & S.W. Wilson (Eds.) *From animals to animats 5* (pp. 13-21). Cambridge, MA: MIT Press.

Keijzer, F.A. (2001). *Representation and Behavior*. Cambridge, MA: MIT Press.

Lyon, P. (2006). The biogenetic approach to cognition. *Cognitive Processing*, 7, 11-29.

Mackie, G.O., Marx, R.M. & Meech, R.W. (2003). Central circuitry in the jellyfish *Aglantha digitale* IV: Pathways coordinating feeding behaviour. *Journal of Experimental Biology*, 206, 2487-2505.

Marder, E. & Calabrese, R.L. (1996). Principles of rhythmic motor pattern generation. *Physiological Reviews*, 76(3), 687-717.

Maturana, H.R. & Varela, F.J. (1980). *Autopoiesis and cognition: The realization of the living*. Dordrecht: Reidel.

Meech, R.W. (1989). The electrophysiology of swimming in the jellyfish *Aglantha digitale*. In P.A.V. Anderson (Ed.), *Evolution of the first nervous systems*. New York: Plenum.

O'Regan, J.K. & Noë, A. (2001). A sensorimotor account of vision and visual consciousness, *Behavioral and Brain Sciences* 24(5), 939-973.

Pfeifer, R. & Scheier, C. (2001). *Understanding intelligence*. Cambridge, MA: MIT Press.

Rowlands, M. (2003). *Externalism: Putting mind and world back together again*. Chesham: Acumen.

Sharkey, N. & Ziemke, T. (2001). Mechanistic versus phenomenological embodiment: Can robot embodiment lead to strong AI? *Journal of cognitive Systems research*, 2, 251-262.

Singla, C.L. (1978). Locomotion and neuromuscular system of *Aglantha digitale*. *Cell and Tissue research*, 188, 317-327.

Sterelny, K. (2001). *The evolution of agency and other essays*. Cambridge: Cambridge University Press.

Stewart, J. (1996). Cognition = Life: Implications for higher-level cognition. *Behavioural Processes*, 35, 311-326.

Van Duijn, M., Keijzer, F.A. & Franken, D. (In press). Principles of minimal cognition: Casting cognition as sensorimotor coordination. *Adaptive Behavior*.

Van Gelder, T. (1998). The dynamical hypothesis in cognitive science. *Behavioral and Brain Sciences*, 21, 615-628.

Wooldridge, M. & Jennings, N.R. (1995). Intelligent agents: Theory and practice. *Knowledge Engineering Review*, 10(2), 115-152.