

PDF hosted at the Radboud Repository of the Radboud University Nijmegen

The following full text is a postprint version which may differ from the publisher's version.

For additional information about this publication click this link.

<http://hdl.handle.net/2066/62486>

Please be advised that this information was generated on 2021-06-22 and may be subject to change.

Form, function and the matter of experience

W. (Pim) F.G. Haselager

Artificial Intelligence/Cognitive Science, Nijmegen Institute for Cognition and Information, University of Nijmegen, Montessorilaan 3, 6525 HR, Nijmegen, THE NETHERLANDS; Pos-graduation Cognitive Science and Philosophy of Mind, Philosophy Department, UNESP, Av. Hygino Muzzi Filho 737, Marília, SP, 17525-900, BRAZIL
pimh@nici.kun.nl

© This paper is not for reproduction without permission of the author(s).

ABSTRACT

The notion of Umwelt (e.g. von Uexküll 1936/2001; 1937/2001) has gained currency in the domain of robotics. Do robots have an Umwelt? Are robots to be conceived of as subjects that truly experience their perceiving and acting upon the world? Or are they merely artificially signaling or behaving as if they do?

Initially, it may have seemed that work in robotics, also known as nouvelle AI, has better chances of dealing constructively with the question of subjective experience than traditional AI because of the embodied embeddedness of its systems. This embodied embeddedness, properly self-organized and dynamically coupled, might be enough to ground the flow of information to such an extent that genuine feeling, volition and intentionality would arise.

However, it has been argued (in two recent papers on von Uexküll; Ziemke & Sharkey 2001; Emmeche 2001) that embodied embeddedness, though important, is not sufficient for establishing the presence of an Umwelt. Systems, it is claimed, need to be alive in order to be able to have an Umwelt. Autopoiesis, the self-producing and self-maintaining property of cells, is characteristic of living systems and is considered to be indispensable in order for a system to have an Umwelt.

I will argue that the relation, that is claimed to exist, between being alive and having an Umwelt is not obvious and in need of considerable clarification. Moreover, I will suggest that the focus on autopoiesis, interpreted as a sharpening of constraints on the matter of implementations can be seen not as in opposition to, but rather as a consequence of the aims of nouvelle AI and its views on the function, form and matter of robotic systems.

1 INTRODUCTION

Throughout the history of Artificial Intelligence (AI)¹, proud programmers that created a program that did something interestingly cognitive have been asked the question; but does it know what it is doing? Does it know that it is doing something? Turing (1950) discussed this question under the heading ‘the argument from consciousness’, and attributed it to Professor Jefferson who spoke about the importance of doing something because of thoughts and emotions felt, and not merely artificially signaled. Recently, it seems that this ‘perennial problem’ of AI has come to plague the nouvelle AI (situated robotics or autonomous agents research) through the notion of Umwelt (von Uexküll 1936/2001; 1937/2001). Are robots to be conceived of as subjects that truly experience their perceiving and acting upon the world? Or are they merely artificially signaling or behaving as if they do?

At least sometimes some robots seem to move around with a purpose, they seem to avoid difficulties and they seem to be capable of sustaining themselves, avoiding damage and energy depletion. They seem to know about their environment and they seem capable of learning more. Some of their reactions to events seem to be based on their history of interactions with the environment. The question is: Is all this mere seeming?

Of course, much is in the eye of the beholder. We, human beings, have a strong tendency to attribute the possession of purpose or volition, thought and beliefs and desires, and even feelings to many things (including, at times, cars and refrigerators) that upon consideration would not qualify as genuinely possessing these capacities. In the case of robots the danger of over-interpretation is present even more strongly (Braitenberg, 1984 gives some amusing examples). This human tendency to over-interpretation may provide the same wind in the sails of scientists working within nouvelle AI as it did for those working within Good Old Fashioned AI (think of Eliza or MYCIN). The potential for commercial exploitation of this tendency is already investigated by companies that build ‘household pet-robots’ (e.g. Sony’s Aibo and SDR-4X).

However, the risk for overinterpretation seems to me to be no greater than the risk for ‘underinterpretation’. For instance, it would serve no useful purpose to exclude, from the outset, the possibility that robots might have the capacity to have experiences. I find arguments such as ‘only living creatures have feelings, purposes and beliefs, robots are not alive, hence they do not have these properties’ to be far from convincing, specifically because the first premise is not well established but rather an issue of empirical investigation. Of course, it may turn out to be true that robots do not have feelings or purposes and beliefs, precisely because they are not living organisms. It is just that I do not think it is valid to reason like this from the outset.

¹ The label ‘AI’, in this context, includes what Haugeland calls Good Old Fashioned Artificial Intelligence (GOFAI: the symbol-rules approach) and traditional forms of connectionism (e.g. the feedforward networks that were particularly popular in the 80s). This form of AI has to be distinguished from more recent work in robotics, that is sometimes addressed as Autonomous Agents Research. However, since the question can and has been raised whether these autonomous agents are genuinely autonomous, I will use the phrase ‘nouvelle AI’ as a general label for recent work in robotics.

The central question I will be concerned with is whether robots have or can have an Umwelt. The notion of Umwelt was introduced in the work of von Uexküll (a.o. 1936/2001; 1937/2001) and it designated the subjective experience of an organism of its perceptual and effector world. It focuses specifically on the phenomenal aspects of specific, perceptually and motorically selected parts of the environment (Emmeche 2001: 3). I think the notion of Umwelt is particularly relevant to nouvelle AI because it emphasizes the *interaction* of the ‘I am interacting with the world’ experience. That is, it stresses more than just the ‘I’, and allows for an approach to experience that does not focus exclusively on the inner aspect of experience.

I will argue that taking the notion of ‘life’ as a necessary condition for the existence of an experienced Umwelt does not help significantly in assessing the capacities of robots. I will suggest that taking a closer look at the way form, function and matter interact may be a more fruitful way to discuss the Umwelt of robots.

2 THE PERENNIAL PROBLEM OF AI

Throughout time, human beings have compared themselves to a great variety of machines. The value of such comparisons has been doubted from the start as well. Hippocrates (around 400 BC), for instance, said the following:

“Comparing humans with their products is an expression of an extraordinary impoverished view on humanity.” (quoted in Simmen 1968: 7-8)

In more recent times, the product with which humans were compared was the clock (cf. Draaisma, 1986). Hobbes (1588 - 1679) raised the question about exactly what properties to attribute to clocks and watches:

“Seeing life is but a motion of limbs, the beginning whereof is in some principal part within; why may we not say, that all automata (engines that move themselves by springs and wheels as doth a watch) have an artificial life” (quoted in Flew 1964: 115).

Descartes (1596-1650) related this question specifically to the debate about animals:

“I cannot share the opinion of Montaigne and others who attribute understanding or thought to animals. (..) I know that animals do many things better than we do, but this does not surprise me. It can even be used to prove they act naturally and mechanically, like a clock which tells the time better than our judgment does. Doubtless when the swallows come in spring, they operate like clocks” (Descartes, 23 November, 1646, letter to the Marquee of Newcastle; Kenny 1970: 206-207).

Based on the same comparison between clocks and organisms, Descartes opposed the suggestion of Hobbes. For Hobbes the self-moving quality of clocks led to the question about whether one could not attribute clocks the property of life, whereas for Descartes the similarity in certain aspects (most notably regularity) of the behavior of clocks and

animals provided sufficient reason to deny animals any form of understanding. De Malebranche (1638-1715) denied that animals experienced anything:

“Animals are without reason or consciousness in the usual sense. They eat without appetite, they scream without pain, they grow without understanding this, they don’t want anything, they don’t fear anything, they are not aware of anything”(quoted in de Wit 1982: 389).

The way he continues is quite interesting in the current context:

“If sometimes perhaps they behave in such a way that this seems reasonable, then this the consequence of a bodily plan, that God ordered, that, on behalf of self-preservation, they without reason, purely mechanically, escape everything that threatens to destroy them” (quoted in de Wit 1982: 389).

If one would replace the word ‘God’ with ‘human being’ and ‘animal’ with ‘computer’ or ‘robot’ a statement results that can be found in the present day in relation to the computational models and robots of AI. Turing, as is well known, discussed the perennial problem of AI under the heading ‘the argument from consciousness’, and attributed it to Professor Jefferson:

“No mechanism could feel (and not merely artificially signal, an easy contrivance) pleasure at its successes, grief when its valves fuse, be warmed by flattery, be made miserable by its mistakes, be charmed by sex, be angry or depressed when it cannot get what it wants.” (Turing 1950: 42).

Basically, we have here the position of De Malebranche, applied to computers instead of to animals.

The perennial problem has been raised in different forms, a.o. ‘they’ (whether clocks, computers, robots, or even animals) are not autonomous, they don’t know what their representations are about, they are not intentional systems, they are not capable of semiosis, they have no originality, they are not creative, they do not have emotions, they have no feelings, they are not conscious, they have no awareness, they are not alive. Something that can confuse the debate considerably is that with whatever question the debate may start, pretty soon one gets tangled up with some of the other questions. Yet, underlying all these questions is the same unifying doubt: is anyone there?² [2]

3 NOUVELLE AI

Robots are interesting candidates about which to raise these questions. Several of their properties seem to make an outright negative answer difficult. First of all, robots are embodied and embedded creatures. That is, they have a body (in contrast with the computational models of traditional AI) through which they interact with their environment (consisting of objects and other (artificial and/or living) creatures, and their embeddedness in the world shapes their behavior and cognitive processes.

² I am ignoring here issues like ‘how would we know whether robots have experiences or not?’ because the same question could be raised in relation to the experiences of other people than oneself (the well known ‘other minds’ problem).

Moreover, much of the behavior of robots seems to be not predetermined but, on the contrary, emergent. Emergence, of course, is a rather murky concept, but in the current context the following aspects are relevant. First of all, emergence in the context of robots can be understood as unprogrammed functionality (Clark 2001: 114). The behavior of the robot is not directly controlled or programmed for in a straightforward way, but arises out of the interactions between a limited number of components that can be substantially different in their properties and action possibilities. Clark gives the example of simple behavioral dispositions (tend towards the right, bounce back when touching something) in a robot that, under the right circumstances, could lead to emergent behavior such as wall following.

Secondly, an important aspect of emergence is that the overall, or higher-level, pattern shapes, influences or constrains the behavior and interactions of the lower-level components. This is sometimes referred to as ‘downward causation’. There have been many debates about how the notion of downward causation should be interpreted in order to make any sense (e.g. Kim 1993). I concur with the position of El-Hani & Emmeche (2000: 262) who claim that downward causality can be understood as a form of (Aristotelian) formal causality:

“Higher level entities establish a particular pattern of constraints on the relations of the lower-level entities composing them.”

The downward causative force of a higher-level patterns can be viewed as restricting the possibilities for interaction among the lower-level components. Finally, the phenomenon of causal spread can be observed in relation to robots. Causal spread is defined by Wheeler & Clark (1999: 106) as follows:

“The phenomenon of interest turns out to depend, in unexpected ways, on causal factors external to the system.”

Traditional AI focuses on what happens *inside* the system. Specifically the central nervous system (artificial or biological) is seen as holding the main causes for the behavior. But according to Wheeler & Clark, the causes of my behavior are not to be found exclusively inside me but they are ‘spread out’ into the environment.

In all, in order to understand the behavior of robots it is necessary to take into account diverse and varying aspects of their body and their environment and the way these aspects interact and self-organize. Thus, it is not unreasonable to investigate the possibility that robots may have an Umwelt.

4 AUTONOMY, UMWELT AND LIFE

In a recent paper, Ziemke & Sharkey (2001: 725-726, 730) examine the Umwelt and autonomy (in the sense of being to a considerable extent independent of their human creators) of robots. Specifically, they focus on robots that evolve through genetic algorithms and that are controlled by recurrent networks. According to them, such robots adapt to their environment and have a historical basis of reaction. That is, the reactions of robots are subjective because they are self-organizing and not completely built in, and because they are specific to the robots and their history of experience. Moreover, robots

are involved in sign processes and make use of the signs themselves, which provides them with a degree of epistemic autonomy. As Ziemke & Sharkey say, robots are ‘on their own’ while interacting with the environment. Third, the development or evolution of robot controllers (i.e. the artificial neural networks) and sometimes even their bodies (in simulation) follows, what von Uexküll called, ‘centrifugal’ principles. They develop from the *inside out*, contrary to the more standard centripetal principle of connecting prearranged parts (like a robot-arm or an optical sensor) to a central unit, from the ‘outside in’. Finally, robots can co-evolve with other evolving entities. Ziemke & Sharkey give a.o. the example of the work by Nolfi & Floreano (1998) where robots (kheperas) controlled by recurrent neural networks co-evolve with other robots into groups displaying either predator or prey behavior. Cliff & Miller (1996) provide an example of internal co-evolution where the controller and the optical sensor evolve in an intermingled fashion.

In all then, one would think that there are good reasons to suspect that robots such as these qualify to a certain extent for having autonomy and an Umwelt. That is, one can give grounds for claiming that, in a rudimentary way, the robots do things on their own and need to have a higher-order mapping and evaluation of their sensory and motor samples of their environment. However, Ziemke & Sharkey end their paper with a clear ‘No’ to the question whether robots such as these have an Umwelt, specifically because these robots are not alive:

“The components might be better integrated after having self-organized, they might even be considered ‘more autonomous’ for that reason, but they certainly do not become alive in that process.” (Ziemke & Sharkey 2001: 736)

The same verdict is given by Emmeche:

“what gives the Umwelt its phenomenal character is not the functional-cybernetic aspect of signal-processing within the system (and at the system-environment interface), but the fact that the living organism is beforehand constituted as an active subject with some agency. Thus, only genuine living beings (organisms and especially animals) can be said to live experientially in an Umwelt.” (Emmeche 2001: 19).

Thus, robots have no Umwelt because they are not alive to start with and they do not become alive in their increasingly autonomous interaction with the world. This argument, if sound, would disqualify artificial creatures with one stroke, and would necessitate robotics to become a branch of biology in order to get any closer to producing creatures with an Umwelt.

At this point, however, I would like to raise a question that may sound strange at first (at least it did to me when I first thought of it): What’s life got to do with it? First of all, ‘life’ and ‘experience’ are not synonyms. The question whether there can be experience without life is an empirical one. Likewise, the question whether artificial creatures (to be distinguished from living creatures) can have an Umwelt is an empirical question. The whole point of research in robotics is to investigate the capacities and properties of robots. It might be a matter of discovery that, due to emergent effects of the couplings between control systems (brains), bodies and environments, the experiencing of an Umwelt may arise in certain kinds of creatures, living or artificial. This is not to say that there can

or will be no differences between artificial and living creatures, but just that the having of experiences need not be a difference.

Secondly, there are situations where creatures can be said to be alive without having experiences. Deep dreamless sleep is normally considered to be experience-less, as are some forms of coma. Organisms without a nervous system are generally considered to be without experience (e.g. Damasio 1999; Emmeche 2001). Hence, being alive is not sufficient for having experiences.

More difficult, of course, is the other issue, regarding creatures having experience without being alive. The suggestion that a non-living creature may have experiences certainly does sound odd. There is a strong tendency to equate 'non-living' with 'dead', and 'being dead' as a state of not (or no longer) experiencing anything. However, it seems that in relation to robots this tendency will not do. Basically, what I am suggesting here is that artificial creatures do not perfectly fit in either the 'dead' or the 'alive' category. Their experiential capacities can therefore not be decided upon by attempts to force them into one of these classes.

So, is life a necessary condition for experience? Well, why would it be? What exactly are the arguments? There are, at least, two arguments to be examined. The first argument is that all experiencing creatures that we know are living creatures. A testimony to the tendency to associate life with experience is that reported after-death experiences are often referred to as 'life after death'. In response one only has to point to the fallibility of induction. It may be the case that up till now all the creatures that can be said to have experiences are living creatures. But this does not constitute proof for the thesis that life is a necessary condition for experience. The unreliability of the inductive argument is all the more pressing because robotics is in the business of constructing creatures of an altogether new (which is why we have reserved the specific word 'robot' to refer to them), artificial, non-biological kind. Exactly what properties apply to these creatures is not something that can be decided upon on the basis of past experiences.

A more substantial argument is one that would explain *why* creatures that can have experiences *must* be alive. One such argument would be that the matter out of which living creatures consist *constitutes* experience. I propose to look more closely at this issue. What properties of living matter would be essential to having experiences?

5 FROM MATTER TO EXPERIENCE: AUTOPOIESIS

I have not been able to find an argument that directly relates the capacity of having experiences to specific properties of matter. One concept, however, that does get mentioned often as being essential to living matter is *autopoiesis* (e.g. Maturana & Varela 1987), referring to the self-generating and self-maintaining capacity of a.o. cells. So, the question becomes: how is autopoiesis related to experience?

Cells are considered to be prime examples of autopoietic systems. However, cells are not generally regarded as having experiences. Emmeche (2001: 18), for instance, indicates that unicellular eukaryotic cells (protozoa) lack a nervous system and do not have a

genuine Umwelt. Thus, it seems that it is not autopoiesis that is constituting the capacity of having experience. Moreover, the notion of autopoiesis does not reflect some intrinsic quality of a specific kind of matter but indicates a characteristic of the *organization* of matter, hence it denotes a functional quality. As Maturana & Varela, (1987: 51; see also Ziemke & Sharkey 2001: 732)

“the phenomena they generate in functioning as autopoietic unities depend on their organization and the way this organization comes about, and not on the physical nature of their components.”

So, autopoiesis could be realized in other stuff. The fact that the particular material that composes living organisms on earth may be a matter of historical or evolutionary coincidence:

“Autopoietic organization can be attained by many different types of components. We have to realize, however, that as regards the molecular origin of terrestrial living beings, only certain molecular species probably possessed the characteristics required for autopoietic unities, thus initiating the structural history to which we ourselves belong” (Maturana & Varela 1987: 49).

If Umwelt, the subjective experience of the perceptual and motor interaction with the environment, turns out to depend on certain *functional* characteristics that are normally found in some forms of matter, why could these functional characteristics not be replicated in other materials? Far from disqualifying robots, this argument leads us to consider more carefully the relation between the function, form and matter of artificial creatures.

6 FUNCTION, FORM AND MATTER

Traditionally, at the center of cognitive science is the idea that any system should be analyzed at three different levels. At the task level, the function of the system is examined, and hypotheses are formulated about what the system has to do. At the formal or computational level, the way in which this function is performed is studied. At this level, it is the form of the processes that is the topic of investigation. And finally, at the physical level, one considers the material composition of the system that is being analyzed. I will, at times, use the terms ‘function’, ‘form’ and ‘matter’ to refer to the three levels of analysis.

From the perspective of traditional AI, the picture that emerges from considering the three levels of analysis is roughly the following. At the task level one considers what the function is that the (sub)system has to fulfill. The analysis at the task level results in a formulation of the main problem that the system has to solve through internal information processing. For instance, Marr (1982) famously stated that the task to be solved by the visual system was the transformation of a 2-dimensional retinal image into a 3-dimensional world of recognized objects. At the organizational level, the focus is on the computational relations between representations.

Algorithms specify the way in which problems can be solved. Different algorithms can be devised and compared in respect to their efficiency and/or psychological plausibility, etc. At the physical level, for traditional AI almost anything qualifies because algorithms can be performed by a whole range of material constructions, from high tech computers to empty beer cans, as Searle has been fond to point out. The analyses at the task

and formal level are such that not too much constraints on the physical level result. Any material substrate that could implement formal processes is, in principle, sufficient.

For nouvelle AI, the situation is quite different. At the task level, the system is thought to be aimed at homeostasis, attempting to maintain its own organization while interacting with the environment. In relation to the form of the organization of the system, the way its components interact, the attention is directed at the way the system self-organizes, allowing a dynamic coupling between system and environment. The requirements of homeostasis and self-organization result in much stronger constraints on the matter in which these processes are to be achieved, resulting in the requirement for autopoietic units. The following matrix summarizes the differences between traditional and nouvelle AI:

	Traditional AI	Nouvelle AI
Function (Task level)	problem solving	Homeostasis
Form (Organizational level)	algorithm	self-organization
Matter (Physical level)	almost anything	autopoietic units

I suggest that nouvelle AI forces us to reconsider the answers provided by traditional cognitive science at the three levels of analysis. Robotics is important to cognitive science in part because the basic ideas presented by nouvelle AI at the three different levels result in much stronger constraints on the *relations* between the levels. This, in turn, leads to a more careful consideration of the material constitution of robots. Importantly, this does not just apply to the artificial nervous systems of robots, but more generally, to the material constitution of the robot as a whole.

7 THE PROTO-UMWELT ARTIFICIAL CREATURES

Let us, by way of conclusion, *consider* what would follow if robots would be developed according to the ideas indicated above. Imagine an artificial system that consists of autopoietic units that maintain themselves. They develop through centrifugal principles, from the inside out. Their coupled processes result in the self-organized emergence of the behavioral patterns that direct the system towards homeostasis. What would we say about systems like these? Perhaps we would not say they were alive, but would they have an Umwelt? Emmeche, for one, gives the following answer:

“If such a material device as a robot could have that special organic flexibility of an animal that allows it to instantiate anything like the law of mind, that is, the tendency to let signs influence or affiliate with other signs in a self-organizing manner, it is difficult to see why such devices should not be able to realize genuine signs (including qualisigns). (...) If that artificially constructed system realizes qualisign action and living feeling, it would have mind in Peirce’s sense. But would it have an Umwelt? Is there anybody home in there, experiencing something?” (Emmeche 2001: 21).

There seems to be something slightly unjust in repeating the question after granting so much. When there is ‘qualisign action’ and ‘living feeling’ it would seem, there is an Umwelt. At some point the perennial question should be stopped, if it is to mean more

than the a repetition of the ‘other minds’ problem. At the same time, one can, perhaps, sympathize with the reluctance to ascribe *full-blown* form of experience to a simple robot.

Maybe we can avoid repeating the perennial question of AI in this situation by considering an intermediate stage on the road to a complete Umwelt. Lucia Santaella, during this workshop, spoke of ‘rudimentary organisms’ like plants and rivers. Perhaps, it would be useful to consider whether robots could qualify as some kind of rudimentary organism having something that perhaps might be called a ‘*proto-Umwelt*’. Interestingly, biologists sometimes speak about cells in a way that comes close to this. Emmeche, for instance speaks about protozoa:

“which lack a nervous system and a genuine Umwelt, but do (..) have a simpler ‘autokinetic or self-moving circle’ by which they enter into semiotic interactions with their exterior milieu.” (Emmeche 2001: 18).

Kaufmann discusses the semiotic capacities of bacteria and notices a similar ‘half-way’ possibility:

“A bacterium knows how to make a living in its world, but no consciousness needs to be attributed. The know-how is the propagating organization. Once there is an autonomous agent, there is a semantics from its privileged point of view. The incoming molecule is yuck or yum. Once yuck and yum, we are not far from Peirce’s meaning laded semiotic triad: sign, signified, significans.” (Kaufmann 2000: 111; also 113).

Finally, the neurophysiologist Damasio (1999: 136-137) indicates that for instance amoeba can be seen as possessing ‘some biological antecedents of the sense of self’. Thus, there is acknowledgment of a form of autokinesis that does lead to something close to semiotic capacities and a sense of self without creating an Umwelt in the full sense of the word. Once we accept this possibility, it becomes hard to see why self-organizing, homeostatic artificial systems that consist of autopoietic units engaging in simple autokinesis and semiotic interaction with their environment would not have at least a ‘proto-Umwelt’ of the kind that is granted to unicellular creatures. Why not attribute such non-living creatures the same kind of ‘proto-Umwelt’ as biologists do to cells?

Acknowledgments

I would like to thank FAPESP for financial support, UNESP, Marilia, SP, Brazil for providing great conditions to work, NICI, Nijmegen, The Netherlands for permission to work at UNESP, João Queiroz and Ricardo Gudwin for a wonderful workshop and Maria Eunice Quilici Gonzalez for comments, warnings, laughter and support.

REFERENCES

- Emmeche, C. 2001. Does a robot have an Umwelt? Reflections on the qualitative bio-semiotics of Jakob von Uexküll. *Semiotica*, 134, 1. 653-693.
- Braitenberg, V. 1984. *Vehicles: Experiments in synthetic psychology*. Cambridge, MA: MIT Press.
- Clark, A. 2001. *Mindware: an introduction to the philosophy of cognitive science*. Oxford: Oxford University Press.
- Cliff, D. & Miller, G. 1996. Co-evolution of pursuit and evasion II: Simulation methods and results, in P. Maes, M. Mataric, J.A. Meyer, J. Pollack & S. Wilson, eds., *From animals to animats 4: Proceedings of the fourth international conference on simulation of adaptive behavior (SAB96)*, pp. 506-515. Cambridge, MA: MIT Press.
- Damasio, A. 1999. *The feeling of what happens: Body, emotion and the making of consciousness*. London: Vintage.
- Draaisma, D. 1986. De machine achter het gordijn: Drie eeuwen debat over bewustzijn en mechanisme in P. Hagoort & R. Maessen, eds., *Geest, computer, kunst*. Amsterdam: Grafiet.
- El-Hani, C. & Emmeche, C. 2000. On some theoretical grounds for an organism-centered biology: Property emergence, supervenience, and downward causation. *Theory of Bioscience*, 119, 234-275.
- Flew, A. 1964. *Mind, body, and death*. New York: MacMillan Publishing Company.
- Kaufmann, S. 2000. *Investigations*. Oxford: Oxford University Press.
- Kenny, A. 1970. *Descartes: Philosophical letters*. Oxford: Basil Blackwell.
- Kim, J. 1993. *Supervenience and mind: selected philosophical essays*. Cambridge: Cambridge University Press.
- Marr, D. 1982. *Vision*. San Francisco: Freeman.
- Maturana, H.R. & Varela, F.J. 1987. *The tree of knowledge: The biological roots of human understanding*. Boston: Shambhala.
- Nolfi, S. & Floreano, D. 1998. Co-evolving predator and prey robots: Do 'arm races' arise in artificial evolution? *Artificial Life*, 4, 311-335.
- Simmen, R. 1968. *Mens & Machine: Teksten en documenten over automaten, androïden en robots*. Amsterdam: van Lindonk.
- Turing, A.M. 1950. Computing machinery and intelligence. *Mind*, 59, 433-460.
- Uexküll, J. von. 1936/2001. An introduction to Umwelt. *Semiotica*, 134, 107-110.

- Uexküll, J. von. 1937/2001. The new concept of Umwelt: A link between science and the humanities. *Semiotica*, 134, 111-123.
- Wheeler, M. & Clark, A. 1999. Genic representation: Reconciling content and causal complexity. *British Journal for the Philosophy of Science*, 50, 103-135.
- Wit, H.C.D. de. 1982. *Ontwikkelingsgeschiedenis van de biologie, Deel 1*. Wageningen: Pudoc.
- Ziemke, T. & Sharkey, N. 2001. A stroll through the worlds of robots and animals: Applying Jakob von Uexküll's theory of meaning to adaptive robots and artificial life. *Semiotica*, 134, 701-746.