Title
Causal Relationships and Relationships between Levels: The Modes of Description Perspective

Journal
Proceedings of the Annual Meeting of the Cognitive Science Society, 21(0)

Authors
Bakker, Bram
Dulk, Paul den

Publication Date
1999

Peer reviewed

Causal Relationships and Relationships between Levels: The Modes of Description Perspective

Bram Bakker (bbakker@fsw.leidenuniv.nl)
Unit of Experimental and Theoretical Psychology; Leiden University
P.O. Box 9555; 2300 RB, Leiden, The Netherlands

Paul den Dulk (dendulk@psy.uva.nl)
Psychonomics department; University of Amsterdam
Roetersstraat 15; 1018 WB, Amsterdam, The Netherlands

Abstract

Many researchers have argued for a description of nature using multiple levels, or modes of description, as we call them. This paper focuses on a confusion that follows from the multiple-mode approach, a confusion due to the notion of causation between modes. Causation between modes is reinterpreted as ordinary causation but with cause and effect described in different modes. In the first part of the paper the framework of modes of description is presented. In the second part it is applied to examples from cognitive science, which are taken from debates on the mind-brain issue and the dynamical systems approach to cognition.

Introduction

Many researchers have argued for scientific analysis at different levels when studying complex phenomena (e.g. Pylyshyn, 1984; Marr, 1982; Newell, 1990; Arbib, 1989; Hofstadter, 1979; Kelso, 1995; Dennett, 1987; Salthe, 1985; Bechtel and Richardson, 1993). Multiple levels—or *modes of description*, as we call them—may be useful to provide a richer picture of a complex phenomenon, and the framework of levels may help thinking about how different descriptions of the same phenomenon relate to each other.

However, the main focus in this paper is on confusions that sometimes arise within this framework, confusions due to the inappropriately mixing up of different modes or levels of description. In particular, we focus on confusions involving *causal relationships* within a multi-level approach.

The next section introduces the framework of modes of description. It is based on a few simple principles, such as the explicit distinction between reality and knowledge and a view of knowledge as many descriptions of a single reality. Our purpose is not to introduce a new, complete philosophical theory of knowledge or to provide a comprehensive overview and rebuttal of all alternative views, but rather to shape and restructure existing (although sometimes disputed) ideas into a simple and coherent framework that is applicable to issues in modern empirical science. Its usefulness when applied to those issues should be the test of its value.

For this reason, we apply the framework to examples of ongoing scientific debates in section 3. The examples concern the mind-brain issue and the dynamical systems approach to cognition.

Theory

One Reality, Multiple Descriptions

Central to the arguments presented in this paper is the straightforward notion that reality and knowledge are different things. Scientific knowledge can be thought of as *descriptions of* reality, but it is not reality itself. Pragmatists (e.g. James, 1907) have argued that reference to reality is unfruitful and unnecessary, but we feel that it is actually more pragmatic, at least for the purposes of this paper, to explicitly assume the existence of a single reality. We hope the usefulness of this assumption will become clear in the remainder of this paper. Thus, we assume that there is only one reality; but many different descriptions of this single reality are possible.

The importance of using multiple descriptions is commonly accepted in cognitive science (see Pylyshyn, 1984; Marr, 1982; Newell, 1990; Arbib, 1989; Hofstadter, 1979; Kelso, 1995; Dennett, 1987). Marr's (1982) implementational, algorithmic, and computational levels of analysis, and Dennett's (1987) physical, design, and intentional stance are particularly well-known examples. In practice, a distinction is often made between the functional level and the implementational level. "Classical", functionalist, symbolic models are viewed by many as descriptions at the abstract, functional level, while connectionist models are viewed as descriptions at a more hardware-oriented implementational level. Neurophysiologists zoom in to low-level physical processes even more when they study chemical processes within a single neuron.

In some cases theories at different levels can be compared as competing theories. For instance, a symbolic model and a connectionist model can compete for the best predictions on the same psychological experiment. They are in competition because the assumption of one reality implies that if two theories make contradictory predictions of data, they cannot both be right. Furthermore, descriptions that explain and predict more data than other descriptions are better descriptions of reality, even if neither of them predicts perfectly.

For such comparisons between different descriptions to make sense, the descriptions must be competing for the same turf, in the sense that they aim to predict the same data. In contrast, very often different levels are intended to predict different data. Functionalist models answer questions about overt behavior, but none about the neural substrate. Neurophysiological models answer questions about the effect of neurotransmitters, but they usually do not predict overt behavior.

Finally, sometimes multiple descriptions need to be considered at the same time in order to make good predictions of data. In that case, they can be said to complement each other (this is an important topic in the paragraph about dynamical systems).

Modes of Description

Some authors identify a level of description more or less directly with physical size or scale (e.g. Salthe, 1985; Yates, 1993). However, sometimes descriptions cannot be neatly arranged within such a hierarchy, and one description cannot easily be thought of as being at a higher level or a lower level than the other. A neuron can be described focusing on electrical properties, or it can be described focusing on chemical processes. Which one is the higher level description? The notion of size is even less of a help when we give a description of the function of a neuron. The "electrical", "chemical" and "function" descriptions are just *different approaches* to understanding reality, different ways of looking at the neuron, all aimed at roughly the same scale. The reason for maintaining them all is that each has its own value for answering different questions.

Conversely, a single approach to understanding reality may be used to describe things at very different scales. Molecules, billiard balls, and planets can often be viewed simply as point masses and be described accordingly, using Newtonian mechanics. The scales are very different, but the description is very much the same and the predictions that follow from this single description are very accurate.

The notion of "levels" suggests a clear-cut hierarchy in which descriptions have an unambiguously defined position that is higher or lower relative to other descriptions. In practice, such a neat ordering seems difficult if not impossible to attain, and we see no reason to search for it. The important point is that there may be different descriptions of the same phenomenon, sometimes differing in scale, sometimes differing in abstractness, sometimes differing in what aspects of the phenomenon the focus is on. This is a more generic outlook on different descriptions than the notion of level can convey, and for that reason we prefer to speak of different *modes of description*.

Concepts

If a theoretical concept plays a pivotal role in a successful description—successful in the sense that it explains and predicts a lot of data—and without this concept the description falls apart completely, we can assume some sort of concurrent *regularity* in reality—a "real pattern", Dennett might say (Dennett, 1991b). All this means is that apparently reality is such that it can be aptly described using that concept. Put the other way around, if there were no such concurrent real regularity, the description could not possibly be as successful as it is. For instance, the concept of atom must have some concurrent regularity in reality, simply because predictions made with theories in which atoms are indispensable are usually extremely accurate. However, arguing that there must be a concurrent regularity in reality is not the same as arguing that there is an *absolute thing* in reality corresponding to the concept in a straightforward way. The exact nature of the regularity cannot be further specified. After all, we only have descriptions, featuring those very concepts, to tell us what reality is like.

Because there are better and worse descriptions, there are better and worse concepts. The concept of atom is better than the concept of phlogiston, because the theory involving atoms is better than the theory involving phlogiston. *Apparently* the concept of atom better corresponds to regularities in reality than the concept of phlogiston does. At the other end, of course, we should also simply discard a concept when it turns out that its corresponding theory is not successful at all, rather than trying to look for the "true meaning" of the concept. After it was discovered that many substances gain weight after combustion rather than lose weight, it was briefly suggested that phlogiston has negative weight (Chalmers, 1982), in an ad hoc attempt to save a concept that had completely lost its value.

Translation of Modes

Even though different modes of descriptions are *separate* ways of looking at reality, they are not *free-floating* ways of looking at reality. Because they are all about one reality, at the very least they cannot contradict each other, in terms of prediction of data. In that sense, different modes must be compatible. What is more, it is often possible to "map" concepts in one mode directly onto concepts in another mode. For instance, the concept of gas—as used in the so-called thermodynamic mode of description—can be translated relatively straightforwardly into concepts used in a molecular mode, when we describe gas as a large number of molecules interacting in a particular way. A molecule (molecular mode), in turn, can often be viewed as a particular threedimensional configuration of atoms (atomic mode).

Sometimes, however, concepts in one mode cannot be translated as easily to concepts in other modes. Consider the concept of life. We cannot map life straightforwardly onto concepts in an atomic or molecular mode of description. If we use the atomic mode of description, we see that atoms and structures of atoms making up living things are basically the same as those making up non-living things. This does not logically imply we have overlooked an important thing or property in the atomic mode that is somehow responsible for life. Nor does it imply that in a life-mode of description something is added to reality that is absent in the atomic mode—we assume there is only one reality. A new mode does not suddenly add things to reality in a mystical way; it only describes this single reality differently. The concept of life in the life-mode corresponds to collective properties and mechanisms of large groups of atoms that are just *not easily described* in the atomic mode. Even though properties described in the life-mode must somehow be realized by mechanisms that can be described in the atom-mode, there is no simple one-to-one mapping from the concept of life to single concepts in the atom-mode. In that sense, not every concept can be "reduced" straightforwardly to concepts in the atomic mode. And for this reason, not everything can or should be predicted from the atom-mode.

Causation

Can we say, at least, that life is *caused* by the (interaction of) atoms? No, we are talking about multiple descriptions—in which the concepts of life and atoms have roles—of one single reality. Different descriptions of the same thing at the same time cannot be said to "cause" each other. Atoms do not cause life, nor do they cause molecules or planets; and life does not cause atoms.

Causation in this sense of levels causing each other tends to get confused with causation in the regular sense, as we shall see in the Applications section. Causation in the regular sense refers to *processes over time, with causes and effects that are not just different descriptions of the same thing.* Suppose we are looking at one neuron in a brain, and the neuron fires. What is the cause of the effect "the neuron fires"?[1] One valid answer is that the presynaptic neurons fired, thus causing our neuron to fire as well. Another valid answer to the same question is that the neuron's membrane potential increased beyond a threshold value, causing the voltage-gated sodium channels to open, causing a massive influx of sodium, leading to a fast and large increase in potential: a spike. Still another valid answer is that an object is present in the visual field that this neuron is sensitive to, so the neuron becomes active.

The first causal story identifies "presynaptic neurons firing" as the cause, and thus stays in the same neural network mode of description as the effect ("the neuron fires"). The other two causal stories identify causes described in very different modes from the neural network mode in which the effect is described: "the membrane potential threshold was crossed" and "the object was present", respectively. Thus, causal explanations may involve *mode switches*: cause and effect need not be described in the same mode. But what makes these examples regular causal explanations is that they all refer to a process over time, and the cause and the effect are not different descriptions of the same thing. In these examples, causation in the other, inappropriate sense would be: "the fast and large increase in potential causes the spike", or: "the feature detection by the neuron causes it to fire" The spike and the fast, large increase in potential are the same thing at the same time, but described in different modes. Likewise, the neuron firing and the detection of the feature are the same thing, but described in different modes.

[1] This example is based on Hofstadter (1981, p. 193–196).

Applications

Mind and Brain

The relationship between mind and brain is one of the classical issues of cognitive science. Descartes (1662) argued that mind and brain are two fundamentally different substances, and although this dualism has lost a lot of ground, there are still many who adhere to it (e.g. Popper & Eccles, 1977; Koestler, 1967). At the opposite end of the spectrum, eliminative materialists argue that the mind is really the brain (Hobbes, 1651, is an early example), and that the mind is at best a fair but limited abstraction of the workings of the brain (e.g. Churchland, 1986). As Armstrong (1968) puts it, "The mind is nothing but the brain.... We can give a complete account of man in purely physio-chemical terms".

Within the modes of description perspective, the crucial step is to consider talk of the mind and talk of the brain as two different modes of description. On the one hand, our earlier assumption of one reality precludes *a priori* a Cartesian duality of mind and brain in reality. On the other hand, this does not mean that we can or need to give "a complete account of man in purely physio-chemical terms". First of all, *all* modes of description are "at best fair but limited abstractions" of reality: the brain mode is still just a description, and not intrinsically "real" (as opposed to the mind mode) because it is "lower" (see Russell, 1962).

The brain mode of description talks about regions in the brain, spiking patterns in networks of neurons, and motor neurons activating muscles. It is used when, for instance, the question is: "How does a person coordinate a hitting movement?". The mind mode of description is a description of reality from "the intentional stance" (Dennett, 1987), and it involves concepts such as intentions, beliefs, desires, the self, and free will. It is used—for now, for a long time to come, and possibly for ever—for answering questions of the kind: "Why does this person hit that other person?". We cannot and need not answer all questions about human behavior using just the brain-mode (or just the mind-mode). Both modes are useful for answering *different* sets of questions, so they can coexist.

Given that both the brain and the mind are valuable modes, what is their relationship to each other? A frequently encountered intuition is that the brain somehow causes the mind. The mind may then be conceived as some sort of "epiphenomenon" of the brain (e.g. Jackson, 1982), an effect of the workings of the brain, but without any real causal role to play in the person's behavior: the mind does not cause muscle fibers to contract, the brain does. That view implies one-way causation from brain to mind. Others—so-called emergent-interactionists—suggest two-way causation: the brain and the mind interact, so that while brain activities give rise to the mind, the mind has "emergent top-down control" over the brain (Sperry, 1991, p. 221). However, just as in the earlier example of atoms and life, the brain mode and the mind mode are two ways of describing one reality, so they do not cause each other, neither one-way nor two-way. Neurons firing do not cause the belief, neurons firing and the belief are the same thing, but described in a different mode.

If relationships between modes are not clearly distinguished from causal relationships, confusions arise. Libet (1985) performed a well-known series of experiments in which he investigated the question whether neural activity caused mental activity or the other way around. Although his research received much criticism, most of it focused on methodological issues. From the modes of description perspective, the question should not even have been asked (see also MacKay, 1985).

Another type of confusion can be found in Hauser (1997). He concludes from his experiments on animal communication (ibid., p. 586):

> These experiments suggest that physiological changes may largely dictate the conditions for call production and suppression, and that an intentional decision need not enter into the equation.

This suggests that the causes for call production and suppression must be mental or physical or possibly a combination of both, and that these options refer to different processes in reality. But because the mental and the physical are just different descriptions of the same phenomenon, talking about deciding between these options makes no sense; an intentional decision in the mind mode must correspond to some physiological changes in the brain mode (even if we don't know what they are).

As a final example of confusion, Sperry (1991, p. 226) vehemently denies Cartesian dualism, but his talk of "emergent interaction" and "downward causation" in mind and brain has led many to reject him as a dualist (e.g. Dennett, 1991a; Vandervert, 1991) or to conclude that even established neuroscience suggests that dualism is true (e.g. Popper & Eccles, 1977, p. 209; Zimbardo, McDermott, Jansz, & Metaal, 1995, p. 90).

If the brain does not cause the mind and the mind does not cause the brain, does this mean that the *mental* event of believing to see a familiar face is not caused by appropriate *physical* stimulation of the retina? Does it mean that a person cannot be said to contract her *physical* muscles because of her own *mental* voluntary decision? No, it doesn't: a causal explanation may contain mode switches, as long as cause and effect do not refer to the same thing at the same time. In this case, description of the different events in the causal chain may be done best in different modes. Stimulation of the retina is an event described in the brain mode which causes, later in time, an event described in the mind mode: believing to see a familiar face. A voluntary decision is an event described in the mind mode, and it causes an event later in time described in the brain mode: motor neurons making muscle fibers contract.

Thus, causal explanations involving voluntary decisions are allowed within the modes of description perspective. Keeping in mind our emphasis on compatibility between modes, that notion may seem incompatible with the basic insight of *determinism*[2] in the physical sciences, but it is not. Concepts from different modes may be *very* different, and the meaning of a concept in one mode does not directly "transfer" to another mode. The notion of voluntary decisions works well in the mind mode of description: people choose freely whether or not to move a limb. In the brain mode or the atom mode there is no such thing as personal choice. But that is a normal state of affairs; in the atom mode there is no such thing as life either. The voluntary decision, described in the mind mode, corresponds to deterministic processes in the brain mode, just as lifelike properties in the life mode correspond to lifeless properties in the atom mode. In other words, demanding compatibility between modes does not mean that concepts from different modes are necessarily mapped onto each other *straightforwardly*. We do not have to be able to identify beliefs with clear-cut brain states in some obvious way in order for us to accept beliefs as good descriptions of cognitive functioning (see Fodor, 1975; Dennett, 1991b). Similarly, we do not have to find voluntary decisions somewhere in the concepts of the brain mode in order for us to accept voluntary decisions as good descriptions of cognitive functioning. Just as in the case of life and atoms, concepts in the mind mode may correspond to collective properties and mechanisms of large groups of neurons—or even the whole person, if we talk about broad concepts such as the self and voluntary decisions (Dennett, 1991a)—that are not easily described from within the brain mode.

The Dynamical Systems Approach to Cognition

One of the most promising and exciting new perspectives in cognitive science is the dynamical systems (DS) approach. It challenges the still dominant idea that the best abstraction of cognitive systems is in terms of "classical", discrete computation and distinct functional modules. Instead of static modules, symbols, propositional logic, and rule-based reasoning, it puts forward the language of dynamical systems, featuring concepts like continuous state spaces, attractors, and bifurcations (e.g. Port & Van Gelder, 1995; Kelso, 1995; Haken, 1983, 1995; Thelen & Smith, 1994; Beer, 1995).

However, there are some confusions in the interpretations of DS models that distract from the important contributions and core issues in the debate. In the DS approach, great emphasis is put on how, in a process called "self-organization", a complex pattern that can be described using the "order parameter" can "emerge" spontaneously when simple units interact. This is a valid and fascinating point, showing that a great deal of regularity and functionality can arise in another, "cheaper", and possibly more robust way than through explicit, detailed instructions that code for the pattern (see Kauffman, 1995).

[2]For argument's sake, we ignore quantum mechanics and deterministic chaos. Deterministic should be read here as "subject to the impersonal laws of nature", being the opposite of "controlled by one's free will".

The relationship between the emerged pattern and the simple units is often referred to as "circular causality" (e.g. Haken, 1995; Kelso, 1995; Keijzer, 1997). Kelso (1995) puts it like this (p. 8–9):

> ...the order parameter is created by the cooperation of the individual parts of the system... Conversely, it governs or constrains the behavior of the individual parts. This is a strange kind of circular causality (which is the chicken and which is the egg?), but we will see that it is typical of all self-organizing systems.

From the modes of description perspective, emergence means that a new mode of description becomes applicable (applicable in the sense that it has predictive and explanatory power) to reality where before it was not applicable, and the order parameter is a concept within this new mode of description. It is basically the same situation as the molecular mode becoming applicable in certain cases, where before only the atom mode was applicable. The order parameter description can be said to *summarize* the behavior of the individual parts (Haken, 1995, p. 26):

> While a huge anount of information is required if we have to describe the behavior of the system by means of the variables q_j of the components, an enormous information compression is achieved by means of the order parameter.

This is one of many cases where a single system ("the system") is described using different modes: one mode focusing in detail on the components q_j, and the other mode providing a compressed description. Therefore, at any one point in time, the emerged pattern cannot be said to be caused by the individual parts or vice versa.

However, causal explanations of the behavior of the individual parts or the order parameter over time may contain two mode switches: change in the order parameter is caused by change in the individual parts, and change in the individual parts in turn is caused by change in the order parameter. In that sense, the total causal story could be called "circular". But it should be clear that this, however interesting, is not a truly new type of causality.

If circular causality is perceived as a truly new type of causality it may give rise to questions such as "how are the causal powers divided between the levels?". It is this type of question that leads to problems, as we saw in the discussion of mind and brain. Kelso sees a strong contrast between circular causality and traditional causality (ibid., p. 9):

> What we have here is one of the main conceptual differences between the circularly causal underpinnings of pattern formation in nonequilibrium systems and the linear causality that underlies most of modern physiology and psychology, with its inputs and outputs, stimuli and responses. Some might argue that the concept of feedback closes the loop, as it were, between input and output. This works fine in simple systems that have only two parts to be joined, each of which affects the other. But add a few more parts interlaced together and very quickly it becomes impossible to treat the system in terms of feedback circuits.

This passage may be read in two ways. If the difference between circular causality and "the linear causality that underlies most of modern physiology and psychology" boils down to the difference between "relationships between modes" on the one hand and "regular causation" on the other, then a meaningless comparison is made between these two entirely different types of relationships. If, alternatively, circular causality is different only because the circularly causal story happens to involve two mode switches, then there is no real "conceptual difference": circular causality and "linear" causality are both just regular causality.

In either case, there seems to be no reason to choose between circular causality on the one hand and concepts such as "inputs" and "outputs" on the other. All systems have input and output. Moreover, there seems to be no reason to discuss the problems of "feedback", because those problems, if any, apply to all systems, including the self-organizing ones.

The unwarranted idea of causation between modes is similarly suggested by phrases like "interaction between levels" (Kelso, 1995; Keijzer, 1997; Salthe, 1985), "leaky levels" (Saunders, Kolen, & Pollack, 1994), "organization across levels" (Keijzer, 1997), "interference between levels" (Salthe, 1985; Keijzer, 1997), and "transactions between levels" (Salthe, 1985). It has led some authors to describe a contrast between systems following a "Newtonian trajectory" and systems following a "regular, biological trajectory" (Yates, 1993; Keijzer, 1997), the latter of which allegedly moves up and down the level hierarchy.

Explaining his idea of levels, Salthe (1985, p.47) says:

> We have seen that current interpretations of complexity tend to lead into matters of scale... we might get into the idea of scale by considering maps. Maps can be drawn to different scales, arbitrarily chosen. A large-scale map might show all of North America... we now decrease the scale of the map, so that we have only the New England coastline showing on a map of the same size...

A map is as clear an example of a description as one can think of: each level is construed as a separate description at a different scale. But later he talks about how lower levels and higher levels "constrain" a certain level of interest, the "focal" level (ibid., p. 82–85):

> Like lower-level constraints, they [higher-level constraints] inform and influence focal-level processes without participating in them dynamically. Where do they come from? In general, from the environment of a process... For some diverse concrete examples: *in vitro* synthesis of sugars and amino acids results in a mixture of racemates while *in vivo* synthesis results only in dextro sugars and levo amino acids; sex determination in turtles depends on environmental temperature; the particular stages a fluid system goes through on its way to turbulence depends upon the width, etc., of the observation chamber.

Here Salthe is talking about how differences in a system's environment cause differences in the system's behavior: causal relationships. Causal relationships and relationships between modes (levels) are explicitly confused. This confusion leads to many peculiar conclusions, such as the conclusion that the interaction between levels must be "limited", because "interactional complexity, of course, would destroy neatly organized level structures" (ibid., p. 53; see also ibid., p.74; Keijzer, 1997, p. 22, 210).

Conclusion

In this paper we have argued against the notion of causation between levels or modes. That notion prompts the question whether a mode is cause, effect, or is interacting with other modes. Actual examples from cognitive science were given to illustrate how this results in confusions and mistaken conclusions. In the framework of modes of descriptions, both cause and effect can be described in many different modes; the issue whether a causal relationship is from one mode to another, the other way around, or interactive, does not even arise. The remaining question is: "In which—possibly different—modes can cause and effect best be described?".

References

Arbib, M.A. (1989). *The Metaphorical Brain II: Neural Networks and Beyond*. New York: John Wiley & Sons.

Armstrong, D.M. (1968). *A Materialist Theory of the Mind.* London: Routledge & K. Paul.

Beer, R.D. (1995). A Dynamical Systems Perspective on Agent-Environment Interaction. *Artificial Intelligence, 72*, 173–215.

Bechtel, W., and Richardson, R.C. (1993). *Discovering Complexity: Decomposition and Localization as Strategies in Scientific Research.* Princeton, NJ: Princeton University Press.

Chalmers, A.F. (1982). *What is this Thing Called Science?* Milton Keynes: Open University Press.

Churchland, P.S. (1986). *Neurophilosophy: Toward a Unified Science of Mind-Brain.* Cambridge, MA: MIT Press.

Dennett, D.C. (1987). *The Intentional Stance.* Cambridge, MA: MIT Press.

Dennett, D.C. (1991a). *Consciousness Explained.* Boston: Little, Brown.

Dennett, D.C. (1991b). Real patterns. *Journal of Philosophy, 87*, 27–51.

Descartes, R. (1662). *The Treatise on Man*, translation from French and commentary by Thomas S. Hall, Cambridge, MA: Harvard Univ. Press, 1972.

Fodor, J.A. (1975). *The Language of Thought.* New York: Thomas Y. Crowell.

Haken, H. (1983). *Synergetics, An Introduction.* Berlin: Springer.

Haken, H. (1995). Some basic concepts of synergetics with respect to multistability in perception, phase transitions and formation of meaning. In: Kruse and Stadler (Eds.) *Ambiguity in Mind and Nature.* Berlin: Springer Verlag.

Hauser, M.D. (1995). *The Evolution of Communication.* Cambridge, MA: MIT Press.

Hobbes, T. (1651). *Leviathan.* London: Crooke.

Hofstadter, D. (1979). *Godel, Escher, Bach.* London: Harvester Press.

Hofstadter, D. (1981) Reflections on: Ant Fugue. In: D. Hofstadter en D.C. Dennett (Eds.) *The Mind's I.* New York: Basic Books.

Jackson, F. (1982). Epiphenomenal qualia. *Philosophical Quarterly, 32*, 127–136.

James, W. (1907). *Pragmatism: A New Name for Some Old Ways of Thinking.* New York: Longmans, Green and Company.

Kauffman, S.A. (1995). *At Home in the Universe.* London: Viking.

Keijzer, F.A. (1997). *The Generation of Behavior.* PhD Thesis, Dept. of Psychology, Leiden University.

Kelso, J.A.S. (1995). *Dynamic Patterns.* Cambridge, MA: MIT Press.

Koestler, A. (1967). *The Ghost in the Machine.* New York: Macmillan.

Libet, B. (1985). Unconscious cerebral initiative and the role of conscious will in voluntary action. *Behavioral and Brain Sciences, 8*, 529–566.

MacKay, D.M. (1985). Do we control our brains. *Behavioral and Brain Sciences, 8*, 546–546.

Marr, D. (1982). *Vision.* New York: Freeman.

Maturana H.M., and Varela J.V. (1984). *The Tree of Knowledge*, Bern: Sherz Verlag.

Newell, A. (1990). *Unified theories of cognition.* Cambridge, MA: Harvard University Press.

Popper, K.R., and Eccles, J.C. (1977). *The self and its brain.* New York: Springer International.

Port, R.F. and Van Gelder T. (Eds.) (1995). *Mind as Motion.* Cambridge, MA: MIT Press.

Pylyshyn, Z.W. (1984). *Computation and Cognition.* Cambridge, MA: MIT Press.

Russell, B.A.W. (1962). *The basic writings of Bertrand Russell.* London: Allen and Unwin.

Salthe, S.N. (1985). *Evolving Hierarchical Systems.* New York: Colombia University Press.

Saunders, G., Kolen, J.F. and Pollack, J.B. (1994). The importance of leaky levels for behavior based AI. *Proceedings of the Third International Conference on Simulation of Adaptive Behavior.*

Sperry, R.W. (1991). In defence of mentalism and emergent interaction. *The Journal of Mind and Behavior, 12*, 221–246.

Thelen, E. and Smith, L.B. (1994). *A dynamic systems approach to the development of cognition and action.* Cambridge, MA: MIT Press.

Vandervert, L. (1991). A measurable and testable brain-based emergent interactionism: An alternative to Sperry's mentalist emergent interactionism. *The Journal of Mind and Behavior, 12*, 201–220.

Yates, F.E. (1993). The Logic of Life. In: C.A.R. Boyd and D. Noble (Eds.) *The Challenge of Integrative Physiology.* Oxford: Oxford University Press.

Zimbardo, P., McDermott, M., Jansz, J., and Metaal. N. (1995). *Psychology: A European Text*, London: HarperCollins.