Penultimate draft of

Haselager, W.F.G., Bongers, R.M. & van Rooij, I. (2003). Cognitive science, representations and dynamical systems theory. In W. Tschacher and J-P. Dauwalder (Eds.) *The dynamical systems approach to cognition: Concepts and empirical paradigms based on self-organization, embodiment, and coordination Dynamics. Studies of Nonlinear Phenomena in Life Science - Vol. 10.* (pp. 229-242). Singapore: World Scientific.

Title

Cognitive science, representations and dynamical systems theory

Authors

W.F.G. (Pim) Haselager, Raoul M. Bongers, & Iris van Rooij

Address

W.F.G. Haselager
NICI, University of Nijmegen,
PO Box 9104,
6500 HE Nijmegen,
The Netherlands
w.haselager@nici.kun.nl

Abstract

In this paper we point out that the assumption of representation in the explanations and models of cognitive science has several disadvantages. We propose that the dynamical systems theory approach, emphasizing the embodied embedded nature of cognition, might provide an important, non-representational alternative. We stress the importance of the challenge, raised by Andy Clark (Clark & Toribio, 1994; Clark, 1997), to dynamical systems theory to deal with 'representation-hungry' cognitive tasks. We indicate a possible way to answer that challenge in a empirically applicable manner. We suggest that investigations of this kind strengthen a motto that can be used as an antidote to the traditional representational cravings of cognitive science: 'Don't use representations in explanation and modeling unless it is absolutely necessary.'

Introduction

The assumption that human behavior is primarily based on the computational processing of representations is undoubtedly part of the very foundation of cognitive science. Over the years, however, certain problems with this assumption have come to light. In this paper we will briefly present some recent developments that, in our view, provide ample reason to reconsider the representational presupposition even in so-called 'representation hungry' cases of higher cognition. Those considerations lead us to believe that an alternative road to understanding those representation hungry cases should be taken. We present an experiment and dynamical model to illustrate the empirical viability of a non-representational approach to cognition.

Representational Headaches

The interpretation of representation we will be focusing on is the one often cited by Haugeland (1991; cf. Clark, 1997a; Clark & Toribio, 1994; Van Gelder, 1995):

> "A sophisticated system (organism) designed (evolved) to maximize some end (e.g., survival) must in general adjust its behavior to specific features, structures, or configurations of its environment in ways that could not have been fully prearranged in its design. [...] But if the relevant features are not always present (detectable), then they can, at least in some cases, be represented; that is, something else can stand in for them, with the power to guide behavior in their stead. That which stands in for something else in this way is a *representation*; that which it stands for is its *content*; and its standing in for that content is *representing* it". (Haugeland, 1991, p. 62).

According to this interpretation a representation is an *internal stand-in* for a (external) state, that is *used by* an organism or system to direct its *behavior*. On the traditional view, internal representations (symbolic or distributed) constitute the knowledge of the system and they function as intermediaries in between perception and action. On the basis of representational computational models cognitive science has achieved great successes in explaining many cognitive processes. However, over the years several problems with the representational presupposition have become apparent.

One of the more important disadvantages of representations becomes clear if one considers the infamous *frame problem* (McCarthy & Hayes 1969; Pylyshyn, 1987; Haselager, 1997). For several decades cognitive science has been trying to make systems that display common sense and that can perform everyday tasks. The task here is not to create models that perform complicated tasks at expert level (e.g. beating the world champion at chess, finding out where large oil reservoirs are hidden). It is about creating models that act 'normally' in real-life, that behave the way humans do in e.g. getting a joke, engaging in some small talk over coffee, finding their way around in a city, etc.

The problem is that 'acting normal' on closer inspection seems to rely on an enormous amount of common sense knowledge. One has to know, for instance, that Dutch people like to think of their southern neighbors, the Belgians, as being very stupid (which they are not) or one would not be able to understand a lot of jokes. Similarly one has to know that people do not like to see their partner running off with someone else, or that under normal circumstances people do not enter a dead end street (although there can be exceptions to that fact), etc. From a modeling perspective, people have an *enormous amount* of common sense knowledge and they are, moreover, capable of *using* that knowledge *swiftly* and *efficiently*. Adequately representing all possibly relevant aspects of the situation and all possibly relevant consequences of actions quickly becomes computationally overwhelming. Thus, a system that should be capable of acting intelligently in the world soon is smothered by its own representational resources. The system just gets lost in its own information store. Too many representations result in computational apathy. The attempts to model common sense reasoning on the basis of large amounts of internally represented knowledge so far have failed.

The frame problem has provided a strong incentive to investigate alternative methods of knowledge representation, such as distributed representations in connectionism. Elsewhere, the first author of this paper has claimed that the problem does not get solved this way, but only gets transformed into a different, but equally serious problem (Haselager, 1997; 1999; Haselager & van Rappard, 1998). Since the frame problem still is as far away from a solution as in the days of McCarthy and Hayes a simple question forces itself upon us: Should everyday knowledge be internally *represented* at all? Thus, the frame problem provides one of the main reasons for taking a critical position regarding internal representations.

During the 1980's the problems encountered by the classical approach became more and more visible to everybody. The symbolic representational format was deemed to be too language-orientated, and too brittle to account for human cognition. The connectionist alternative, based on distributed representations, was greeted enthusiastically. Flexible weight-settings and activation-patterns across groups of units replaced the representational building blocks of the 'Language of Thought' (Fodor, 1975). Quite soon, however, problems with distributed representations became clear as well. Neural networks turn out to have great difficulty in representing information that specifies who is doing what with whom and with what (as in: 'John hit Paul with the hammer') in a way that allows it to be used in reasoning. Finding a solution to this problem of 'systematicity' turned out to be possible only in limited domains. Moreover, the development of useful distributed representations is overly dependent on the prior, and by itself rather brittle, arrangement of training sets that have to be run through an astonishing number of times (e.g. Aizawa, 1997; Chalmers, 1993; Churchland, 1987; Fodor & Pylyshyn, 1988; Hadley, 1994; Haselager & van Rappard, 1998; Horgan & Tienson, 1997; Shastri & Ajjanagadde, 1993; Smolensky, 1988).

Of course connectionists do not fail to point out the weaknesses of the classical symbolic representational scheme, whereas the classically inclined relentlessly indicate the difficulties experienced by the connectionists. The important but frustrating thing is: *Both* sides are right. It is a matter of concern, we think, that the foundation of cognitive science, i.e. representation, is subject to fierce debates about what constitutes the right representational format. It gets even more serious when the criticisms of both representational formats are basically *correct*. The question then becomes whether the corner stone of cognitive science really is all that solid after all.

A further point to note is the increasing *omnipresence* of representations. In the 'good old days' of cognitive science, when the behaviorists had just been beaten, everyone knew what representations were: Well structured symbol strings with a well-regulated syntax and semantics. Automated formal systems simulated human cognitive processes that were supposed to take place in a 'Language of Thought' (Fodor, 1975). An important, but often neglected, side-effect of this clarity was that it was also clear which systems were *not* representational. As Fodor (1986) indicated, only systems that can respond selectively to non-nomic properties of a stimulus are candidates for the ascription of mental representations.

We think that it is important for a notion to have genuine content that cases are provided to which the notion does not apply. However, over the years, the notion of representation has become applied to increasingly larger domains. By now, the notion seems to be applicable to every system that can be interpreted as performing a function with at least one intermediate step, such that the "output" of one step serves as "represented input" for the subsequent step. This results in simple physical machines such as the Watt Governor, threshing machines and even sieves qualifying as representational and computational systems (van Gelder, 1995; Bechtel, 1998; Churchland & Sejnowski, 1992; Haselager, de Groot & van Rappard, in press). In our view, such results are cause for concern, because if representations are ubiquitous, the notion threatens to loose its substance. After all, the notion of representation was introduced to account for cases of cognition that behaviorism could not properly deal with. But if representations turn out to be omnipresent, what surplus value can they have for the study of cognition?

Of course, we do not wish to claim that these points are knock-down arguments against the representational assumption of cognitive science. Clearly, the amount of empirical success that has been achieved on a representational basis is impressive and cannot be ignored or downgraded on the basis of certain theoretical considerations, serious as they may be. However, we do feel that the problems outlined above provide enough incentive to investigate whether the representationalist's monopoly in cognitive science can be avoided.

We submit that it is wise to consider once again just which cases necessitate a representational account and which do not. Put differently, it may be a good research strategy to see just how far one can go without assuming representations from the start. Thus, the motto we will be defending is not 'Away with representations', but rather 'Don't use representations in explanation and modeling unless you really have to'.

Embodied Embedded Cognition and Dynamical Systems Theory

The alternative approach to the study of cognition that we are investigating has been discussed widely over the last decade under a variety of labels (e.g. enactive cognition, embodied cognition). We wish to focus specifically on the embodied embeddedness of cognition and contrast this with the idea of internal representations. The basic idea of embodied embedded cognition is that cognition cannot properly be understood without a recognition of the importance of bodily interaction with the environment. This implies a shift in focus; instead of assuming internal representational and computational resources from the start, one sets out by examining the possibilities for action that the body and the environment offer the organism. Likewise, the environment is not seen primarily as raising *problems* for the organism that it has to solve by means of internal information processing, but rather as providing *opportunities* for interaction.

This alternative view implies a different conception of the *task* that the cognitive system is facing: Cognition is not primarily about problem solving and planning, but about discovering and responding to the available possibilities for action. From this perspective, it may be more fruitful to consider the organism not as a store of represented knowledge about the world, but rather as embodying a behavioral repertoire that is being tapped into by the environment.

In all then, the embodied embedded view of cognition claims that a lot of what we know need not consist of represented knowledge about the environment, because for everyday behavior most of the relevant information is out there in the world for us to see and/or use. Moreover, in as far as we 'store' knowledge, it might be more fruitful to interpret this knowledge not as information about the world but rather as embodied behavioral patterns that resonate to the environment in which we are embedded.

It is, of course, a considerable task to make the notions of 'embodiment' and 'embeddedness' precise in a way that allows experimental investigations on the basis of them. Moreover, formal models should be created that can be compared with the traditional representational-computational models of cognitive science. It is for this reason that Dynamical Systems Theory (DST) can be of great value to cognitive science. It provides the formal means for analyzing how the couplings between body, brain and environment leads to behavioral patterns that unfold over time.

Despite the fact that DST has become a well-established research method (e.g. Haken, 1977, 1983, in press; Haken, Kelso, & Bunz, 1985; Kelso, 1995, in press; Thelen, in press; Thelen & Smith, 1994) its impact on cognitive science is a matter of continuing debate. A prevailing view is that though DST is of great relevance to simple, often cyclical, behavior, its importance for understanding 'higher' cognition has yet to be established. Such a position has been expressed in a constructive way by Clark (Clark & Toribio, 1994; Clark, 1997) in the form of a challenge for DST to deal with *representation-hungry* cases that, for example, "include thoughts about temporally or spatially distant events and thought about the potential outcome of imagined actions." (Clark, 1997, p.167).

A Representation-Hungry Task

It is far from obvious how one should go about studying a representation-hungry case in a systematic way; i.e., a way that would afford explicit dynamical modeling. For this reason we chose to start with a simple example of a representation-hungry task; viz., the judgment of whether or not it would be possible to reach a distant object with a hand-held rod (for a more detailed description of this experiment and its results, see van Rooij, Bongers, & Haselager, 2002). This task requires a subject to form a judgement about the possible outcome of a future action. Though simple, the task is representation-hungry according to Clark's definition. From a traditional cognitive science perspective, 'judged reachability' is conceptualised as involving some form of computation over (1) represented distance, (2) represented body dimensions, and, in the case where a tool is used, (3) represented aspects of the tool.

The main aim of our study of judging reachability was to show that behaviour depending on imagination *can* be experimentally investigated and modelled from a dynamical perspective. In our experiment participants were positioned at a predetermined distance from an object placed on a table. On each trial they were handed a rod and where asked to indicate, by simply saying "yes" or "no", whether or not they believed they would be able to reach the object with the rod if they would try to do so. Across trials rod length was manipulated systematically; i.e., either rod length increased in 1.5 cm increments from minimum to maximum (increase sequence: I), or decreased in 1.5 cm increments from maximum to minimum (decrease sequence: D), or the rod lengths ranging from minimum to maximum were randomly assigned to the task (random sequence: R). Further, increase and decrease sequences were always coupled and followed by a random sequence, resulting in IDR and DIR blocks (in these blocks the random sequence functioned as a buffer in between coupled ID or DI runs). Our analysis of the results indicated interesting effects that could be captured by a two-attractor model we adopted from Tuller, Case, Ding, & Kelso (1994).

A Two-Attractor Model

Within the DST approach many different models have been developed to account for global patterns in behavior. Given that the task we studied involved discerning which rods enabled successful reaching and which did not, we used a dynamical model particularly designed to account for behavior with two attractor states. Tuller et al. (1994; see also Case, Tuller, Ding, & Kelso, 1995) applied such a model to speech categorization phenomena. Following the example of Tuller et al. (1994) we use equation 1 to model our data.

$$V(x) = kx - \tfrac{1}{2}x^2 + \tfrac{1}{4}x^4 \tag{1}$$

in which x denotes the set of possible states and k is the control parameter specifying the direction and the degree of tilt of the potential function (c.f. Tuller et al., 1994). The potential function V(x) has maximally two attractors; e.g. when $k = 0$. The attractor for positive x is arbitrarily assumed to correspond to "no" (i.e., the participant indicates the belief or judgment that it is not possible to reach the object with the rod) and the attractor for negative x is assumed to correspond to "yes" (the participant indicates the belief or judgment that it is possible to reach the object with the rod. As can be seen in Figure 1, for $k < -k_c$ and for $k >$

$+k_c$ only one attractor exists, namely “no” or “yes” respectively. This means that for $k < -k_c$ the response will always be “no” and for $k > +k_c$ the response will always be “yes”. For $-k_c < k < +k_c$ the two attractors coexist, and hence both “no” and “yes” are possible judgments. The latter situation is referred to as *multistability*, which in this case specifically means *bi*stability. Given the nature of attractors (i.e., that they *attract* the system to organize accordingly) multistability can lead to *hysteresis*. Hysteresis is the phenomenon that the switch, in this case from “no” to “yes”, occurs for a higher value of the control parameter than the switch from “yes” to “no”. We will illustrate this effect by means of Figure 1.

Insert Figure 1. about here

In Figure 1 the little black dot represents, for each value of k, the state in which the system is most likely to settle if k would systematically *increase*. That is, for $k = -1$ only one attractor exists (i.e., “no”), and the system will settle in it. Increasing k forces the function to tilt, but even though the no-attractor has become more shallow, it is still the only attractor, and hence the ball stays in the no-attractor. When the control parameter reaches the critical value $-k_c$ the yes-attractor starts to appear. However, even though now multistability exists the system will tend to cling to the no-attractor simply because it has already settled in that one (i.e., it takes too much effort to “climb the hill” to get to the yes-attractor). At $+k_c$, however, the attractor corresponding to “no” ceases to exist, and hence the “no” state becomes a repellor and the system unavoidably will settle in the yes-attractor. Now, if the same story is told for systematically *decreasing* k one observes hysteresis, since in that case the system will tend to cling to the yes-attractor throughout the multistable region. Note that random influences can cause the system to switch between the two attractors when multistability exists (which is the reason why we talk in terms of a *tendency* to cling to the present attractor-state). Because such random influences are inherent in any concrete, embodied system the dynamic system can probably best be interpreted as a stochastic system in modeling real-world data.

To capture participants’ behavior in our task the relationship between the control parameter k and the independent variable, in our case rod length, has to be specified. Following Tuller et al. (1994) we assume that this relationship is not a one-to-one correspondence. Instead we view k as a function of (1) rod length, (2) the number of repetitions of the categorical judgments[1], and (3) perceptual and cognitive characteristics of the participant. The relationship between the control parameter and rod length can be symbolized by the equation 2:

$$k = ? + (N_{no} – N_{yes})S, \qquad (2)$$

in which k specifies the value of the control parameter, ? is linearly proportional to the length of the rod, N_{no} and N_{yes} are growing functions of the number of accumulative repetitions of “no” and “yes” respectively, S?0 and represents relevant characteristics of the participant that may fluctuate during the time course of the experiment[2]. Given that S represents uncontrolled factors influencing task behavior, we cannot know the exact value of S. Therefore, we take a qualitative approach to the combined influences of (N_{no}-N_{yes}) and S on the dynamics of the behavior of the participants.

Insert Figure 2. about here

In equation 2 if (N_{no}-N_{yes})S=0 then k=?. So when either N_{no}=N_{yes} or S=0, then there is a one-to-one correspondence between k and ?. However, for S>0, k will be larger than ? when N_{no}>N_{yes} and k will be smaller than ? when N_{no}<N_{yes}. The rationale of Equation 2 is illustrated by Figure 2 for a coupled sequential run, in which rod length first systematically increases and subsequently systematically decreases (an ‘ID-run’). For the sake of clarity we hold S constant and only look at the effect of the accumulative repetitions of a response.

In an ID-run the participant is presented at first with the smallest rod (bottom left in Figure 2). For short rods the participants start with no-responses and N_{no} will become increasingly larger than N_{yes} (which will remain zero) with every next trial. Due to the fact that N_{no} grows increasingly larger than N_{yes}, k will increase faster than ? increases. When k reaches the value of +k_c a transition occurs and the participant switches to yes-responses. With every next trial N_{yes} will grow, whereas N_{no} will not. Hence the slope of the function k will decrease. Because the increase sequence is followed by a decrease sequence N_{yes} will start to outnumber N_{no}. This will cause k to decrease faster than ?. When -k_c is reached a

[1] See also Parducci’s (1965; Parducci & Wedell, 1986) range-frequency theory and Helson’s (1964) adaptation-level theory.

[2] See Tuller et al. (1994) for a slightly different specification of the relationship between the control parameter k and the experimental variable ?.

transition occurs and the participant will switch to no-responses. Figure 2 thus illustrates that for sufficiently large S the transition from “yes” to “no” occurs at a larger rod length than the transition from “no” to “yes”. This is an example of the *enhanced contrast* effect.

Note that in terms of where the switch from “no” to “yes”, and vice versa, occurs on the rod length continuum, “enhanced contrast” is exactly the opposite effect of “hysteresis”. The two-attractor model can account for both effects by assuming that hysteresis is caused at the level of the control parameter, and enhanced contrast occurs because the relationship between control parameter and rod length is modulated by accumulated repetitions. When $(N_{no}-N_{yes})S=0$ then k = ?, and thus hysteresis is observed (see Figure 1). However, when S >> 0 then k grows or shrinks faster than ? and enhanced contrast may be observed (Figure 2). Further, one can imagine that for a certain settings of the parameters the hysteresis and enhanced contrast effects ‘cancel out’ and one finds that the first and second transition occur at exactly the same rod length. This finding is referred to as *critical boundary*. However, the number of parameter settings that result in critical boundary is much smaller than the number of settings that result in either hysteresis or enhanced contrast.

Model Predictions and Results

A couple of predictions can be derived from the interrelationship between Equation 1 and 2 described above. In this context we will derive three predictions and discuss how our experimental findings relate to these predictions. In our original experiment more hypotheses were tested (cf. van Rooij, Bongers & Haselager, 2002) but for brevity we only select a representative set of predictions that suffice to illustrate the qualitative fit between model and data.

First, we said that a dynamic system tends to remain in the state it resides in if that state is an attractor state. In terms of the data obtained in random sequences one would expect this tendency to express itself in an *assimilative bias*. To test this hypothesis we investigated the association between the response on the present trial (“yes” or “no”) and the response on the preceding trial (“yes” or “no”). As expected we found an assimilative bias; i.e., the probability of a “yes”-response (“no”-response) on the present trial was highest when the response on the previous trial had been “yes” (“no”) as well.

Second, we expect that the assimilative bias (as observed in random sequences) leads to hysteresis in coupled sequential runs (i.e., runs in which first rod length systematically increases and then decrease, ID, or vice versa, DI). However, hysteresis may not be observed when the contrastive influence of accumulated repetition is sufficiently large (see Figure 2).

We found that in our experiment enhanced contrast was the most frequent pattern observed in individual coupled sequential runs (±73% of all coupled runs), but also a considerable number of times hysteresis was observed (±21% of all coupled runs). As expected, critical boundary was the least common response pattern, occurring only on ±6% of all coupled runs, and no more than once per participant. We interpret the finding of hysteresis as an indication that, in these cases, the modulation of the relationship between rod length and control parameter (based on accumulated repetitions of judgments) was not sufficiently strong to cause enhanced contrast to occur.

We were able to test this claim, due to the fact that overall the participants tended to overestimate their reaching distance. As a consequence the average 50% point was much lower than the middle of the range of rod lengths employed in the experiment. This led to an exceptionally large number of repetitions of “yes” responses in ID-runs as compared to DI-runs. In accordance with the assumption that the crucial source for enhanced contrast is the modulating effect that accumulation of repetitions have on the control parameter, we found that hysteresis was much more likely to occur in DI-runs than in ID-runs, while the reverse held true for enhanced contrast.

Third, as can be seen in Figure 2, in the two-attractor model the contrastive influences of response-repetitions is naturally coupled with a decrease in the size of the multistable region (i.e., as measured in terms of ?). Since the number of opportunities for switches can be expected to be positively related with the size of the multistable region, a negative relationship between number of repetitions and the probability of observing switches is predicted by the model. In accordance with this prediction, we found that less switches occurred in the second part of coupled sequential runs (e.g., D in ID), where the repetitions accumulated to a greater extent than in the first part (e.g., D in DI). Further, this difference in frequency of switches between first and second part of a coupled run, was most prominent for decrease sequences, as compared to increase sequences. Again, this can be seen as a consequence of the fact that participants tended to overestimate the distance reachable, leading to an extraordinarily large accumulation of repetitions of “yes” in the D-part of ID runs.

Discussion & Conclusion

In sum, we submit that the dynamical model captures the experimental results in a straightforward way. In our experimental task there is a clear coupling between aspects of the body (a.o. reaching distance, stability of position), features of the environment (a.o. the distance to the object, rod length) and cognition (i.e. the judgment based on imagined action). Contrary to traditional expectations, these features need not be represented in order to model the experimental results. Instead, the dynamics of the embodied embedded behavior that we examined could be accounted for by means of a fairly simple model.

Of course, a representational account *might* be able to explain these data as well. One could, for instance, explain hysteresis by means of the residual activation of representations, or enhanced contrast by means of the deactivation or inhibition of representations. However, we find it difficult to see to what extent a representational account could explain the *systematicity* in the occurrence of hysteresis and enhanced contrast. More importantly, we wonder why one even would want to search for a representational model if there is a simple and sufficient dynamical account available. It is the traditional cognitivist desire for a representational account that we want to put into question.

To traditional cognitivists, however, the dynamical model may be intuitively unconvincing. An often heard argument is that a representational account provides a *mechanism* that explains how the data come about whereas a dynamical account is merely a convenient way to *describe* the data. We think, however, that this is a mistake. *Both* the representational and the dynamical models specify the constraints that govern the behavior of the actual mechanism that the body, brain and environment implement. Both models are formally specified, and both can be simulated by means of a computer program. The important difference between the representational and dynamical models lies not in their relation to the actual mechanism, but in what the models presuppose concerning the *nature* of the mechanism. Specifically, it is the absence of representations that makes the dynamical model more economic (more conform Occam's razor) and hence more attractive.

Our study shows that a non-representational account of imagined action may well be possible. This observation is important as it provides incentive for further study of other representation-hungry cases from a dynamical point of view. We should *test* how far we can go without assuming representation, before we assume that representations are indispensable. Hence the motto: Don't use representations, unless they are really necessary.

References

Aizawa, K. (1997). Explaining systematicity. *Mind & Language, 12*, 115-136.

Barsalou, L.W. (1991). Deriving categories to achieve goals. *The psychology of learning and motivation, 27*, 1-64.

Bechtel, W. (1998). Representations and cognitive explanations: Assessing the dynamicist's challenge in cognitive science. *Cognitive Science, 22,* 295-318.

Beek, P.J., Peper, C.E., & Stegeman, D.F. (1995). Dynamical models of movement coordination. *Human Movement Science, 14,* 573-608.

Case, P., Tuller, B., Ding, M., & Kelso, J.A. (1995). Evaluation of a dynamical model of speech perception. *Perception & Psychophysics, 57,* 977-988.

Chalmers, D. J. (1993). Connectionism and compositionality: Why Fodor and Pylyshyn were wrong. *Philosophical Psychology, 6,* 305-319.

Churchland, P.S., & Sejnowski, T. (1992). *The computational brain.* Cambridge, MA: MIT Press.

Clark, A. & Toribio, J. (1994). Doing without representing? *Synthese, 101,* 401-431.

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

Fodor, J. A. (1975). *The language of thought.* Cambridge: Harvard University Press.

Fodor, J.A. (1986). Why paramecia don't have mental representations. In P. French, J.T. Uehling & H. Wettstein (Eds.), *Midwest studies in philosophy* (pp.3-23). Minnesota: University of Minnesota Press.

Hadley, R.F. (1994). Systematicity in connectionist language learning. *Mind & Language, 9,* 247-272.

Haken, H. (1977). *Synergetics: An introduction.* Berlin: Springer Verlag.

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

Haken, H., Kelso, J.A.S., & Bunz, H. (1985). A theoretical model of phase transitions in human hand movements. *Biological Cybernetics, 51,* 347-356.

Haken, H., & Wunderlin, A. (1990). Synergetics and its paradigm of self-organization in biological systems. In H.T.A. Whiting, O.G. Meijer, & P.C.W. van Wieringen (Eds.), *The Natural-Physical Approach to Movement Control*. Amsterdam: University Press.

Haken, H. (in press). Intelligent behavior - A synergetic view. In W. Tschacher & J-P. Dauwalder (Eds.) *Dynamical Systems Approaches to Embodied Cognition.*

Haselager, W.F.G. (1997). *Cognitive Science and folk psychology: The right frame of mind.* London: Sage.

Haselager, W.F.G., de Groot, A., & van Rappard, J.F.H. (in press). Representationalism versus anti-representationalism: A debate for the sake of appearance. *Philosophical Psychology.*

Haselager, W.F.G., & van Rappard, J.F.H. (1998). Connectionism, systematicity and the frame problem. *Minds and Machines, 8*, 161-179.

Haselager (1999). On the potential of non-classical constituency. *Acta Analytica, 22,*23-42

Helson, H. (1964). *Adaptation-level theory: An experimental and systematic approach to behavior.* New York: Harper & Row.

Horgan, T., & Tienson, J. (1997). *Connectionism and the philosophy of psychology.* Cambridge, MA: MIT Press.

Kelso, J.A.S. (1995). *Dynamic Patterns: The self-organization of brain and behavior.* Cambridge, MA: MIT Press.

Kelso, J.A.S. (in press). Cognitive Coordination Dynamics. In W. Tschacher & J-P. Dauwalder (Eds.) *Dynamical Systems Approaches to Embodied Cognition.*

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

McCarthy, J., & Hayes, P.J. (1969). Some philosophical problems from the standpoint of artificial intelligence. In B. Meltzer & D. Michie (Eds.), *Machine intelligence*. Edinburgh: Edinburgh University Press.

Parducci, A. (1965). Category judgment: range-frequency model, *Psychological Review, 72*, 407-418.

Parducci, A., & Wedell, D. H. (1986). The category effect with rating scales: Number of categories, number of stimuli, and method of presentation. *Journal of Experimental Psychology: Human Perception and Performance, 12,* 496-519.

Port, R., & Van Gelder, T. (1995). *Mind as motion: explorations in the dynamics of cognition.* Cambridge, MA: MIT Press.

Pylyshyn, Z.W. (Ed.) (1987). *The robot's dilemma.* Norwood: Ablex.

Shastri, L. & Ajjanagadde, V. (1993). From simple associations to systematic reasoning: a connectionist representation of rules, variables and dynamic bindings using temporal synchrony. *Behavioral and Brain Sciences, 16*(3), 417-494.

Smolensky, P. (1988). On the proper treatment of connectionism. *Behavioral and Brain Sciences, 11,* 1-74.

Thelen, E. (in press). Grounded in the world: Developmental origins of the embodied mind. In W. Tschacher & J-P. Dauwalder (Eds.) *Dynamical Systems Approaches to Embodied Cognition.*

Thelen, E., & Smith, L.B. (1994). *A Dynamic Systems Approach to the Development of Cognition and Action.* Cambridge, MA: MIT Press.

Tuller, B., Case, P. Ding, M., & Kelso, J.A. (1994). The nonlinear dynamics of speech categorization. *Journal of Experimental Psychology: Human Perception and Performance, 20,* 3-16.

Van Gelder, (1992). Defining ‘distributed representation’. *Connection Science, 4*(3&4), 175-191.

Van Gelder, T. (1995). What might cognition be, if not computation? *Journal of Philosophy* , *92,* 345-381.

Van Gelder, T. (1998). The dynamical hypothesis in cognitive science. *Behavioral and Brain Sciences, 21,* 615-665.

Van Rooij, I., Bongers, R.M., & Haselager, W.F.G. (2002). A non-representational approach to imagined action. *Cognitive Science (26)*3, 345-375.