CC AI

THE JOURNAL FOR THE INTEGRATED STUDY OF
ARTIFICIAL INTELLIGENCE
COGNITIVE SCIENCE AND APPLIED EPISTEMOLOGY

CONCEPTS AND COMPLEX SYSTEMS

Philip Van Loocke (ed.)

VOLUME 13 1996 NUMBERS 2-3

ISSN 0773-4182 Quarterly Journal

CC-AI

THE JOURNAL FOR THE INTEGRATED STUDY
OF ARTIFICIAL INTELLIGENCE, COGNITIVE SCIENCE
AND APPLIED EPISTEMOLOGY

Editorial Address: CC-AI
Blandijnberg 2
B-9000 Ghent, Belgium

Tel.: (32)(9)264.39.51
Telex: RUGENT 12.754
Telefax: (32)(9)264.41.97

CC-AI, ISSN 0773-4182, published by COMMUNICATION AND COGNITION
Blandijnberg 2
B-9000 Ghent, Belgium

Cover design, Yves De Vlieger

THE PERSPECTIVE OF SITUATED AND SELF-ORGANIZING COGNITION IN COGNITIVE PSYCHOLOGY

Wolfgang TSCHACHER
University Psychiatric Services
University of Bern
Switzerland

Christian SCHEIER, lic.phil.
AI Lab
Computer Science Department
University of Zurich
Switzerland

Abstract

We discuss a theoretical framework of cognitive psychology that allows for an understanding of the adaptivity, goal directedness, and flexibility of behavior. Goals and intentions as explanatory principles were banned from academic psychology under the influence of behaviorism. With the advent of the information processing view of cognitive psychology, this taboo has been overcome, but scientific understanding of intentionality is still lacking. At present a computational view of cognition and action dominates throughout psychology. Such current syntactical models are usually descriptive and make strong assumptions concerning internal representations; they imply a manipulation of symbols and categories which are supposed to correspond to entities in the world. Other recent theories in cognitive psychology are oriented more toward motivational constructs; they are based on volition and intention as explanations for action regulation. These latter theories therefore encounter the problem of teleology, because they rely on semantic homunculi in the mind which allocate attention, retrieve information from memory stores, and develop intentions, enabling the individual to act.

In our view, two approaches may be helpful to achieve a coherent new theoretical framework for cognitive psychology. First, synergetics and self-organization research provide principles of pattern formation and adaptivity which can be applied to complex systems such as the mind. Second, 'New Artificial Intelligence' (New AI) and the situated cognition approach have criticised classical AI research for being in quite a similar kind of impasse as

cognitive psychology is. Consequently, the approach of 'situated and self-organizing cognition' claims that emergent patterns in cognition regulate action in an adaptive manner. Cognition is situated by control parameters ('valences' which express environmental constraints). Optimality of patterns is achieved by synergetic dynamics in the valence-driven mind.

1. Problems of cognitive psychology and action psychology

The current state of affairs in cognitive psychology is characterized roughly by two approaches.

The first approach is the view of *information processing* which goes back to the 'cognitive turn' psychology took in the 1960s (Miller et al., 1960). At the heart of the information processing view is the notion of computation of mental symbols which represent real-world entities (Fodor, 1975). The mechanisms of cognition are supposed to be implemented as cybernetic control of hierarchical feedback loops (Carver & Scheier, 1982). Motivational variables and intervening variables (like self-efficacy expectations: Bandura, 1977) result from perceived discrepancies of actual state and goal state. Computer science and classical AI have been important driving forces for the introduction of the information processing framework in psychology.

The second approach is that of *action theory*, which has a long-standing tradition in psychology. The concept of action defined as goal-directed and planned behavior is deeply rooted in psychological introspection and philosophy (Aristotle's 'causa finalis'). Early psychological theories of willed action (Ach, 1910; Lewin, 1926) are elaborated in today's volitional psychology (Heckhausen & Kuhl, 1985). Volition research focusses on the cognitive and motivational analysis of intentions as determinants of action control (Heckhausen et al., 1987).

The two approaches rest on opposing premises. Cognitive information processing derives from a technical notion of a computational model of the mind (Anderson, 1983; Langley, 1983). This view is syntactic and cybernetic in nature. In action theory, on the other hand, motivational variables (wishes on the emotional side, intentions on the cognitive side) as a means of self-regulation are primary. This is more compatible with folk psychology in that it assumes behavior to be self-controlled and intentional. Thus, the action approach is basically semantic. Both approaches share the assumption of mental models and mental representations of the world, upon which the mind is supposed to act. Only after the various stages of mental processing have been passed can the stage of realizing action finally be entered.

Both approaches encounter serious problems which have prompted us to look for alternative paths of conceptualization. Which problems are these? As the two approaches described above are diverse, we shall start with discussing each separately. For the sake of clarity, we will not explicitly address the efforts which have been made to combine the two approaches (e.g. Kuhl, 1992, who claims that "... a functional treatment of *self* within a computational theory of mind is possible.").

The problems of the computational view have been discussed extensively by many authors in AI (see below) and in psychology (e.g. Kolers & Smythe, 1984). It has been shown, for example, that at least the initial stages of attentional processes are massively parallel rather than serial (Neisser, 1976). Recognition of patterns is not easily understood as a process reducible to bitwise processing of information. This point has already been made in Gestalt psychology (Wertheimer, 1912; Helson, 1933), which put forward a holistic theory stating that an array of features is not just the sum of the features of all components but a different entity, a 'gestalt'. In other words, there may be emergent properties in perception and, generally, in cognition, which can hardly be accounted for by a computational approach.

In simulations of computational models of cognition several shortcomings become apparent. This may be one of the reasons why most of the criticisms of the computational approach have been formulated in AI, whereas there is no comparable debate in psychology (for exceptions, see e.g. Haken & Stadler, 1990; Thelen & Smith, 1994). For example, one problem which emerged in AI research concerns *learning* in tasks which demand unsupervised learning (see below). More generally, the computational view does not handle *change* and the dynamics of cognition very well. This point is addressed by connectionist AI where learning is studied in the context of neural networks. We will discuss this approach in more detail below.

Additionally, we may also look at the other side of the coin of dynamics, at *stability*. It is vitally important that cognitive entities remain stable under quite diverse conditions. This is evident in perception where objects are perceived as invariant even if they are changed, distorted, occluded, transposed, etc. In social cognition, belief systems and attitudes are maintained in different environments and under different circumstances. Generally, a concept may be applied meaningfully to different sets which have not one component in common.

Concepts, therefore, are not simply linear compositions of elements but must have higher-level qualities which can be evoked in a non-symbolic way. Thus, stability in the face of various transformations is an important attribute of cognition.

In conclusion, it seems that basic *dynamical* attributes of cognition — learning and stability — which are fundamentals of any cognitive processing, are not readily understood by computational theory (see also Vallacher & Kaufman, 1996).

Many of these criticisms do not apply to the action theoretical approach. On the contrary, those phenomena which are hardest to define in computational terms are the premises of action theory. The setting and pursuit of goals by a self-determined agent is prerequisite to the very definition of action. The main problem here is that we are dealing with concepts which are teleological right from the start, so that an old debate arises: how can goals (i.e. *future* states of an individual) cause intentions and wishes which determine the individual's behavior in the *present*? Obviously, this formulation may be an adequate description of anybody's introspection but is not a scientific explanation.

Action theory claims its intentionalistic terminology can explain behavior. This situation resembles pre-Darwinian biology: the giraffe has a long neck because it intends to eat from trees. But we may accept the terminology of action theory as a sort of abbreviated, *descriptive* code for mechanisms which have yet to be explained, non-intentionalistically and in detail. In section 4, we 'intend' to do just this.

Another assumption of most theories of cognitive psychology has recently come under vigorous attack, namely the assumption of mental representations. The computational view of representation is based on the premise that there is a 'language of the mind' (Fodor, 1975); the world is mapped into the mind in a logical (propositional) or analogous (mental models) fashion. Thus the world is represented by mental tokens (categories) and the categories are processed according to computational syntactical rules. But research on categorization and on memory has shown that this is probably not the whole story. Human categories do not have the 'classical' properties of set-theoretical categories where membership is defined by singly necessary and jointly sufficient conditions (Rosch & Lloyd, 1978). As we have already stated above, categories are dynamical entities which are diffuse *and* stable in a way which is appropriate under the given constraints of a social and cultural environment. Furthermore, the rules by which thinking connects categories and concepts are often not

identical with the rules of probability and logic (see, e.g., the conjunction fallacy in decision making, Kahneman et al., 1982). This again shows that computation is not a sufficient explanation of human cognition.

How can these problems be overcome?

Cognitive psychology presents a rationalistic picture of cognition: action is seen as behavior that progresses systematically from wishes to intentions, according to a plan, schema or script. Decision is determined by 'value x expectation' considerations. There are hierarchies of feedback loops. In short, in the foreground of this view we deal with a fixed cognitive architecture applying fixed formal algorithms to symbols which represent the world in an unambiguous manner.

But what we know from observation and self-observation seems quite different: there usually is a flow of thought, ideas, intentions, emotions. All mental events are incessantly changing even in the absence of environmental change (although we are not concerned here with the neurological substrate of cognition, this applies also to brain activity). Creative and adaptive ideas and actions may come out of a broad 'stream of consciousness' which is perceived as being beyond control and planning. We doubt that there is as much a priori structure in our cognitions and actions as the cognitivists tell us.

Therefore we suggest that cognition should be conceptualized differently: we will not take structure for granted but start from the flow of thought. Dynamical pattern formation can serve as an alternative to pre-wired computation. Thus, we can ask the opposite question: how does cognitive architecture emerge from cognitive dynamics?

Several attempts have been made in psychology to investigate alternatives to the predominant computational theory of mind. We have already mentioned gestalt psychology. Gestalt theory gave way to behavioristic theory (i.e. anti-cognitive information processing) in the middle of this century, but specialized species of gestaltlike conceptualizations still exist. One of them dwels in the ecological approach to perception, which was put forward by Gibson (1979) and Kugler & Turvey (1987). Gibson, who had been a student of Lewin's, developed an 'ecological theory of perception'. Its central term is *affordance*, a concept which links ecological stimuli directly to the perceiver. Usually (when enough ambient information is available, i.e. outside the tachistoscopic lab), perception consists of a direct 'pick up' of relevant information.

In philosophy, affordances reflect the property of 'Zuhandenheit' (readiness -to-hand) (Heidegger, 1962/1927). In a 'hermeneutical cycle', perception may be seen as interpretation or as understanding of objects based upon some pre-knowledge, which has its roots in culture and phylogeny. Therefore, all symbol processing is based on a history of antecedents because it is embodied; expressed in terms of the recent debate, symbol processing is always situated.

This point is elaborated in psychology, among others, by Greeno's situativity theory. Situativity is seen as a general characteristic of cognition (Greeno & Moore, 1993; Law, 1993), which puts the focus more on environment-organism coupling than on cognition 'inside the mind'. As this debate — situated action vs. computation — is a core topic of cognitive science, we shall elaborate on it in the next section.

2. The problems of classical AI

There seems to be consensus within a large part of the research community in AI that classical systems are brittle, that they lack integrated learning and generalization capabilities, and that they cannot perform in real time. This makes them ill-suited for real world applications. We will not discuss these problems in detail as they are well-known. We would rather focus on the underlying reasons for these problems because they were instrumental in the New AI approach (Pfeifer and Scheier, submitted).

The frame problem

The frame problem was originally pointed out by McCarthy & Hayes (1969). It has more recently attracted a lot of interest (e.g. Pylyshyn, 1987). The central issue concerns how to model *change* (Janlert, 1987): given a model of a continuously changing environment, how can the model be kept in tune with the real world? Assuming that the model consists of a set of logical propositions (which essentially applies to any representation in classical AI) any proposition can change at any time. For example, consider propositional representations such as: INSIDE(ROBOT, ROOM) or ON(BATTERY, WAGON).

Assume that there is a set of such representations of the environment stored in a robot's memory. There is a battery and a time bomb on a wagon. The task of the robot is to remove the battery from the room and recharge it in a safe place. The problem here is one of determining the implications of an action. For example, the action of moving the wagon has the 'side effect' that the bomb will

also be moved. Unfortunately, the robot does not know that this is relevant. What is entirely obvious to a human observer has to be made explicit for a robot.

The first idea is to have the robot take possible 'side-effects' into account. There are potentially very many. Checking them all takes a lot of time and most are entirely irrelevant. Another solution might be to try to distinguish between relevant and irrelevant inferences. But in order to do this one has to consider them all anyway which implies that this approach does not have a significant advantage over the former one.

The frame problem really is about the system-environment interaction. The question is how models of a changing environment can be kept in tune with the environment. This is not a problem of logic, but rather one of modeling.

In the real world it is not necessary to build a representation of the situation in the first place: one can simply look at it, thereby disburdening oneself of cumbersome updating processes. Moreover, we can point at things when talking about them (see also 'situatedness' below).

The frame problem is a fundamental one and is intrinsic to every modeling approach. As soon as there is a model of a changing environment, there is a frame problem. An important goal of intelligent systems design in New AI is to minimize the implications of the frame problem. One of the ways to achieve this is to minimize the amount of modeling in the first place.

The symbol grounding problem

The symbol grounding problem refers to the question of how symbols relate to the real world. In classical AI the meaning of symbols is typically defined in a purely syntactic way by how symbols relate to other symbols and how they are processed by some interpreter (Newell & Simon, 1976; Quillian, 1968). The relation of the symbols to the outside world is rarely discussed explicitly. In other words, we are dealing with closed systems. This position not only pertains to AI but to computer science in general. Except in real-time applications, the relation of symbols (e.g. in database applications) to the outside world is never elaborated, it is assumed as somehow given, the — typically implicit — assumption that designers and potential users will know what the symbols mean (e.g. the price of a product). Interestingly enough this idea is also predominant in linguistics: it is taken for granted that there is some kind of correspondence between the symbols or sentences and the outside world. The study of meaning then relates to the translation of sentences into some kind of logic-based representation where the semantics is clearly defined (Winograd & Flores, 1986, p. 18). This position is acceptable in the area of natural language since there is always a human interpreter and it can be safely expected that he or she is capable of establishing the appropriate relations to some outside world: the mapping is 'grounded' in the human's experience of his or her interaction with the real world.

However, once we remove the human interpreter from the loop, as in the case of autonomous agents, we have to take into account that the system needs to interact with the environment on its own. Thus, the meaning of the symbols must be grounded in the system's own interaction with the real world. Symbol systems in which symbols only refer to other symbols are not grounded because the connection to the outside world is missing. The symbols have meaning only to a designer or a user, not to the system itself.

It is interesting to note that for a long time the symbol grounding problem did not attract much attention in AI or cognitive science — and it has never been an issue in computer science in general. Only the renewed interest in autonomous robots has pushed it to the foreground. This problem has been discussed in detail by Harnad (1990). It can be argued that the symbol grounding problem is really an artifact of symbolic systems and 'disappears' if a different approach is used.

The problem of situatedness

The concept of situatedness has recently attracted a lot of interest and led to heated debates about the nature of intelligence and the place of symbol processing systems in studying intelligence. For example, a complete issue of the journal *Cognitive Science* is dedicated to 'situatedness' (Cognitive Science *17*, 1993). 'Situatedness' roughly means the following: First, it implies that the world is viewed entirely from the perspective of the agent (not from the observer's perspective — see the 'frame-of-reference' problem below). Second, a situated agent capitalizes on the system-environment interaction. Its behavior is largely based on the current situation rather than on detailed plans. Third, a situated agent brings its own experience to bear on the current situation; depending on its experience, it will behave differently. In other words, it changes over time. As it turns out, situated agents, i.e. agents having the property of situatedness, are much better at performing in real time because while exploiting the system-environment interaction they minimize the amount of central processing.

The perspective of situatedness contrasts with traditional AI where the approach has been — and still is — to equip the agents with detailed models of their environment. These models form the basis for planning processes which in turn are used for deciding on a particular action. But plan-based systems quickly run into combinatorial problems, i.e. the frame problem. If the real world changes, one of the main problems is keeping the models in tune with the environment. Inspection of the problem of taking action in the real world shows that it is neither necessary nor desirable to develop very comprehensive and detailed models (e.g. Suchman, 1987, Winograd & Flores, 1986). The more comprehensive and detailed the models, the harder the agent will be struck by the frame problem.

Typically, only a small part of an agent's environment is relevant for its action. In addition, instead of performing extensive inference operations on internal models or representations, the agent can interact with the current situation. The real world is, in a sense, part of the 'knowledge' the agent requires in order to act (we put 'knowledge' in quotes to indicate that this is not the standard way of using this term in AI. The standard way refers to knowledge structures that are represented internally.). The agent can merely 'look at it' through the sensors.

Traditional AI systems are not situated and there is no reason why they should be for there is always a human interpreter in the loop. However, if we are interested in building (or understanding) systems which act directly in the real world, they must be situated. Otherwise, given the properties of the real world, the system will not be able to perform intelligently, i.e. in real time, taking only the relevant aspects of the situation into account.

The frame-of-reference problem

Whenever we are involved in designing an intelligent system, we have to be aware of the 'frame-of-reference' problem. Our outline of the problem is based on Clancey's extensive treatment (Clancey, 1991). He argued that if we want to build models using computers or robots, we must appropriately conceptualize the relation among the observer, the designer (or the modeler), the artifact, and the environment. This problem is called the 'frame-of-reference' problem.

The first thing we must understand is that behavior is always the result of a system-environment interaction. In order to clarify this point, let us refer to the example of Simon's ant on the beach (Fig. 1; Simon, 1969). Let us assume that an ant starts moving on the right-hand side of a beach and its anthill is somewhere on the left. The direction it travels is roughly from right to left. The path the ant might take will be arduous because the beach is full of pebbles, rocks, puddles and other obstacles. But this complexity may, in fact, be only an apparent one. It would be a frame-of-reference mistake to conclude from the — apparent — complexity of the trajectory that the internal mechanisms which are responsible for generating the behavior of the ant also have to be complex. The mechanisms which drive the ant's behavior may be very simple, implementing 'rules' that we would describe as follows: "if obstacle sensor on left is activated, turn right" (and vice versa). In interaction with the environment, the apparent complexity of the trajectory emerges. The 'rules' are patterns in the neural structures of the ant (how simple patterns may originate in a neural system, which is complex even in the ant, will be discussed in section 4).

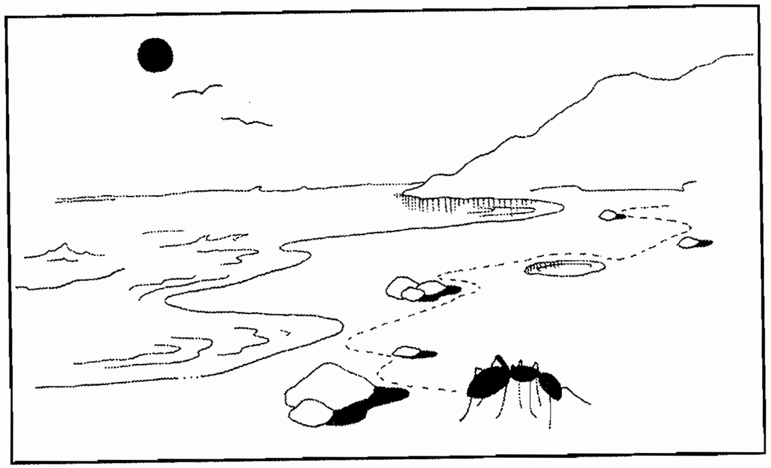

Figure 1: Simon's ant on the beach

Note that the seeming complexity of behavior emerges from the interaction and not from the environment alone: it would be just as erroneous to claim that the complexity of the trajectory is due to the complexity of the environment. The complexity of the environment is only a prerequisite. If we would increase the size of the ant, say, by a factor of 100, and let it start in the same location with exactly the same behavioral rules as before, it would go more or less in a straight line! What appeared to the normal ant as obstacles would no longer be recognized as such by the giant ant. Its antennae would not be sensitive enough to detect the irregularities on the beach.

Or take a human sitting on the beach. Introducing other humans will make the environment much more complex and interesting as this offers the potential for

highly sophisticated interactions (talking, playing, kissing, etc.). Now replace our first human by an ant. To the ant it is entirely irrelevant whether an object in the environment is a human or any other moving object: it does not have the sensors and brain system to experience the complexity.

Starting in the early 1980s, many of the problems of classical AI have been claimed to be solvable using the principles of connectionism. For example, neural networks are not brittle when confronted with new learning input as are propositional systems. They show 'graceful degradation' when information is incomplete or noisy. Moreover, certain classes of learning algorithms, namely the non-supervised learning schemes (e.g. Kohonen feature maps, Hopfield nets) are compatible with the self-organization view advocated later in this paper. They can be shown to form patterns and achieve pattern recognition based on the local interactions of subsymbolic components.

Nevertheless, they do not resolve the more fundamental problems mentioned in this section; neural networks still face the frame-of-reference problem. For example, the activation of output nodes in the standard back-propagation networks must be interpreted by a human observer — thus, there is always an observer in the loop. Also, neural networks *always* learn; they do not (have to) distinguish between relevant and irrelevant input stimuli. All items of the learning sets are relevant because the observer typically has made a careful pre-selection of the input data.

We suggest that neural networks are viable tools if they are embedded in a complete system that interacts autonomously with its environment. Only in this way can neural networks circumpass the frame-of-reference problem. There is no need any longer for an observer to interpret the activations of the nodes in the networks, rather the network is grounded in the physical body and the interactions with the real world (Harnad, 1990). This point, which is at the core of New AI, will be discussed in more detail below.

3. The New AI approach: complete autonomous agents

New AI has emerged as an alternative to classical AI, as an attempt to overcome the problems of the traditional approach. The core idea is to study the interaction of an agent (human, animal or robot) with the environment or real world, rather than investigating well-defined problems in virtual or block worlds. In other words, the focus is on the situated activity of an agent in its environment. Intelligence is seen in the interaction, not within the system.

This approach has several implications. We shall focus on three of them: embodiment, completeness, and ecological niches.

Embodiment

Interaction with an environment implies that *embodied* systems have to be constructed. Embodiment is a prequisite of situated cognition. Only if a system is in direct relation with its environment, i.e. only if it has a body of some sort, is it able to act in a situated way.

In terms of intelligent systems design and modelling the idea of embodiment has led to a remarkable increase of work with *mobile robots* or *autonomous agents.* Mobile robots constitute the optimal tool for New AI, because on the one hand they allow for a synthetic approach (the hallmark of AI), and on the other hand can be used to study issues of situated cognition implemented in interaction with the real world. Let us look at one example.

Brooks' subsumption architecture was the first approach towards New AI, or behavior-based robotics (Brooks, 1986). It is a method of decomposing the control architecture of a robot into a set of task-achieving 'behaviors' or competencies. The usual approach of conceptualizing intelligence is based on functional decomposition: First, there is sensing (i.e. perception), then internal processing (e.g. world modeling, planning, decision making) and finally some actions are executed (e.g. moving forward, grasping an object). This leads to the sense-think-act cycle of the traditional information processing approach. It is sometimes also called horizontal decomposition since each module follows the other sequentially (see Fig. 2). In contrast with the traditional approach, subsumption architecture builds control architectures by incrementally adding task-achieving behaviors on top of each other. Implementations of such behaviors are called layers. Higher level layers (e.g. WANDER) build and rely on lower level ones (e.g. AVOID). Higher layers can subsume lower layers. Hence, instead of having a single sequence of information flow — from perception to world modeling to action — there are multiple paths which are parallelly active. Each of these paths (or layers) is concerned with only a small subtask of the robot's overall task such as avoiding walls, circling around targets, moving to a charging station, etc. These layers can function relatively independently. They do not have to await instructions or results produced by other layers. In short, the subsumption approach realizes the direct couplings between sensors and actuators, with only limited internal processing, and can therefore tackle the frame problem mentioned above.

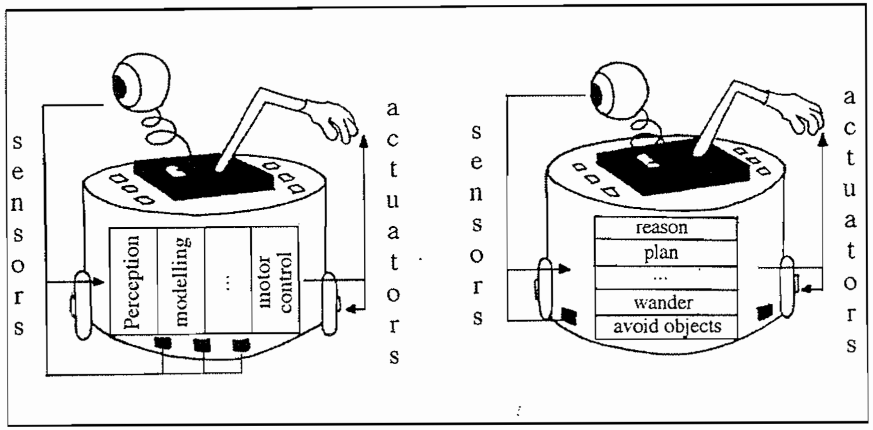

Figure 2, left: The traditional decomposition of a control architecture into a sequence of horizontally layered functional modules. Right: Vertical decomposition based on task-achieving behaviors in the subsumption approach (from Pfeifer & Scheier, p. 134).

Complete Systems

In 1961 the Japanese psychologist Masanao Toda proposed studying 'Fungus Eaters' as an alternative to the traditional methods of academic psychology (Toda, 1982). Rather than performing ever more restricted and well-controlled experiments on isolated faculties (memory, language, learning, perception, emotion, etc.) and narrow tasks (memorizing nonsense syllables, letter perception on degraded stimuli, etc.) we should study 'complete' systems, though perhaps simple ones. This idea is fundamental to the research agenda of New AI. 'Complete' in this context means that the systems are capable of behaving autonomously in an environment without a human intermediary. Such systems have to incorporate capabilities for classification, navigation, object manipulation, and for 'deciding' what to do. The integration of these capabilities into a system which is capable of behaving on its own, so the argument went, will yield more insights into the nature of mind or intelligence than looking at fragments of the unbelievably complex human mind. The 'Fungus Eater' approach can be seen as a precursor of a more ecologically-minded psychology (e.g. Neisser, 1976, 1982; Neisser & Winograd, 1988).

The 'Fungus Eater' is an autonomous agent sent to a distant planet to collect uranium ore. The more ore it collects, the more it will be rewarded. It feeds on a certain type of fungus which grows on this planet. The 'Fungus Eater' has an internal fungus store and means of locomotion (legs), means for decision making (brain) and collection (arms). Any kind of activity, including 'thinking', requires energy. If the level of fungus in the store drops to zero the 'Fungus Eater' is dead. The 'Fungus Eater' is also equipped with sensors, one for vision and one for detecting uranium ore (e.g. a Geiger counter).

The scenario described by Toda is interesting in a number of respects. The 'Fungus Eaters' must be autonomous: they are simply too far away to be controlled remotely. They must be self-sufficient as there are no humans to replace the batteries and to repair the robots, and they must be adaptive because the territory in which they function is largely unknown.

Ecological niches and universality

If we look at biological agents, i.e. animals, we find that they require an environment for survival which is suited to satisfy their needs. Such an environment is called an 'ecological niche'. Wilson (1975) gives the following definition: "the range of each environmental variable such as temperature, humidity, and food items, within which a species can exist and reproduce." (p. 317).

In nature, there is no such thing as a 'universal animal.' Animals (and humans) are always 'designed' by evolution for a particular niche. Agents behave in the real world. As pointed out, they always require certain conditions for their survival. A robot always requires some kind of energy source. It must be equipped with sensors and effectors in order to perform its task in a particular environment, or more precisely, in a particular 'ecological niche'. If the robot has to work at night, it may not be a good idea to equip it only with a vision sensor: an infrared device might be necessary. So, the idea of an ecological niche holds for robots as well. It follows that there can be no universal robot, a constraint deriving from the fact that it has to perform in the real world.

This contrasts sharply to computation. Computation is universal: Turing machines are the only machines that need to be studied. This is, of course, only possible because computation, by definition, 'takes place' in a virtual world. And universality only applies in this virtual world. Computers are sometimes said to be universal. This is true only when focusing on computation. If we look at computers as being real machines, they depend very much on their environment. They require a supply of electricity, must be handled by their users with care, must not be exposed to excessive heat, etc. In this sense, computers, just like

any other artifact, are designed for a particular ecological niche. Of course, some robots can perform several tasks and exist in more varied environments compared to others, so their niche is broader, but nonetheless still there.

The fact that agents in the real world are not universal but have to function in a particular niche, sounds like a severe restriction. However there is a lot of leverage to be gained, too. Because the ecological niche is restricted, has its own laws and characteristics, its types of objects and agents, its temperature profile (i.e. how temperature changes over time), its lighting conditions, etc., there is no need to provide for everything in the agent itself. Assume that in a particular niche only large objects are relevant. Then there is no need for a high-resolution sensor for distinguishing really small objects. If the niche is flat, wheels are sufficient. Often, learning problems which at the purely computational level seem intractable, converge in real-time if the constraints of the econiche are exploited. For example, it might be the case that all objects of interest have a bilateral symmetry, which implies that learning can be unilateral. This makes life much easier. However, as always, there is a tradeoff: the more constraints we exploit in our designs, the less universal the agent will be.

The goal of New AI is to build autonomous, self-sufficient, situated, embodied agents designed for a particular ecological niche. We shall not go into any further detail here because this would be far beyond of the scope of this paper (for reviews of New AI research, see Pfeifer & Scheier, submitted; Steels & Brooks, 1995).

4. Synopsis: situated and self-organizing cognition

In this last section, we will outline a framework for a cognitive 'architecture' that dispenses with an architect, but still has the capacity to account for organized and rational action. In the theoretical framework to be presented below, cognition and goal-oriented action are viewed as emergent properties of a self-organized cognitive system.

Which phenomena have to be addressed by such a framework ? Let us first — as in a 'Gedankenexperiment' — picture the mind as a bundle of innumerable rudimentary cognitive items, a set of fleeting cognitive and emotional micro-events, a 'stream of consciousness'. In terms of dynamics, we address a system spanning a very highdimensional phase space. Daringly, we may name this bundle a 'cognitive-emotional system', CES for short. There would initially be no coherence and pre-wired structure in this system, just the microscopic chaos of the local behavior of the cognitive components (to be defined later). We might find 'associations', i.e. some local interaction between components which happen to be in some (temporal or spatial) vicinity (all spatial terms are to be understood based on the notion of phase space of dynamical systems theory, i.e. as an abstract space which is usually not the Euclidean space; Abraham & Shaw, 1992).

We would have to expect this system's behavior to be very complex, with almost as many degrees of freedom as the number of items or dimensions it consists of. Obviously, such a system is not a good choice for a model of the mind. What would be lacking?

1) Pattern — A CES would have to devise a way in which its components can be structured and organized. Components must be grouped, sorted, combined, etc. on a large scale, depending on their relevance for different tasks and demands. Thus, if some situation requires all the cognitive items that pertain to, say, writing an essay, then other items adequate for repairing a bicycle should be relatively less active. Of course, an outside designer / interpreter or an internal homunculus is not allowed to provide for pattern.

2) Stability — As soon as some cognitive-emotional pattern is established it should be stable over time, random or irrelevant changes in the environment and in the CES must not result in immediate restructuring. The stable state of a dynamical system ('attractor') may be defined as a state whose neighbors in phase space remain in the former's vicinity. Asymptotic (global) stability means that (all) perturbations to an attractor are damped out with time. As long as a pattern within the CES is stable, the activity of 'writing an essay' is not (at least not necessarily) transformed into 'repairing a bicycle' should, for example, a bicycle happen to pass by the window of my study.

3) Optimality — Writing essays or repairing bicycles may temporarily be inadequate actions for an individual. Thus, not only should a CES have the potential to form patterns with some asymptotic stability, but the patterns and their stability should also be useful in a given situation. There must be a function that 'tells' the CES how to shape its patterns in order for it to be optimally adapted to an environment.

Self-organization

We propose that an answer to these demands may lie in the phenomenon of self-organization which is modeled comprehensively by synergetics (Haken, 1983). Self-organization is ubiquitous in complex systems which are in a far-from-equilibrium state (Prigogine & Stengers, 1984). Archetypal physical systems like the Bénard instability and the laser are well-known examples of such systems; these systems are capable of pattern formation not imposed by an external designer.

The Bénard instability can be demonstrated in a layer of fluid (Kratky, 1992). When the fluid is heated from below while the medium at its upper surface is kept at a constant lower temperature, a heat flux permeates the fluid system. As soon as the heat flux reaches a critical value, highly coordinated patterns of convection appear in the fluid. In addition to local interactions among fluid molecules, a much stronger long-range interaction is imposed. Patterns take the shape of rolls or of hexagonal cells, depending on the form of the container and the site of the heat application (Fig. 3). Upon a further increase of heat flux (i.e. of the difference in temperature above and below the fluid) patterns change first to oscillating rolls and later to nonperiodic patterns showing deterministic chaos. The Bénard instability and equivalent atmospheric systems can be modelled quite well by equations of only three degrees of freedom (Lorenz, 1963).

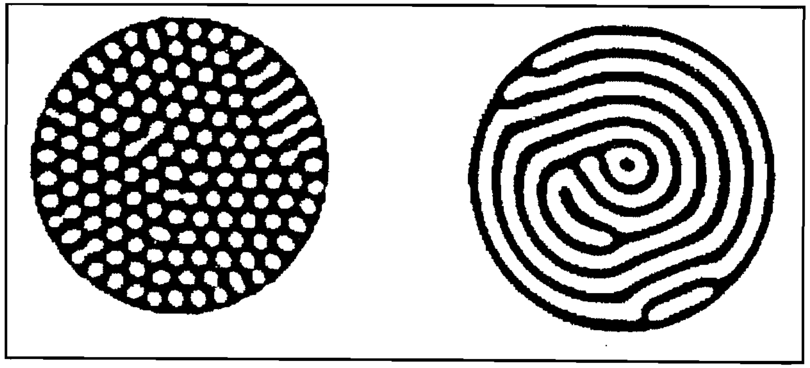

Figure 3: Bénard hexagonal and roll patterns seen from above (after Bestehorn et al., 1993)

In summary, what is remarkable about this behavior is the emergence of a highly ordered macroscopic pattern out of the random microscopic movements of fluid molecules, i.e. of many degrees of freedom. Furthermore, the capacity of the system to transport heat is increased when its control parameter grows; loosely speaking, the system adapts to its non-equilibrium environment by 'trying to reduce' the gradient of temperature. All three attributes listed above — pattern formation, stability and optimality — can be found in the dynamics of the Bénard instability and other self-organizing natural systems.

A general model of self-organizing systems is illustrated in Fig. 4. This model has three constituents: control parameters are variables of the system's environment which denote the system's departure from thermodynamic equilibrium (i.e. the temperature gradient in case of the Bénard instability). Control parameters 'drive' the complex system (the particles of the fluid); after a phase transition this driving produces patterns (quantified by order parameters). Thus the complex system (endowed with many degrees of freedom at the start) has become a two-level system: it may now be described completely at a macroscopic level by the few degrees of freedom of order parameters (in our example, by specifying the regular convection patterns). The two levels of the system are linked recursively; order parameters emerge from the microscopic dynamics, and in turn structure ('enslave') the motion of the system's many microscopic components.

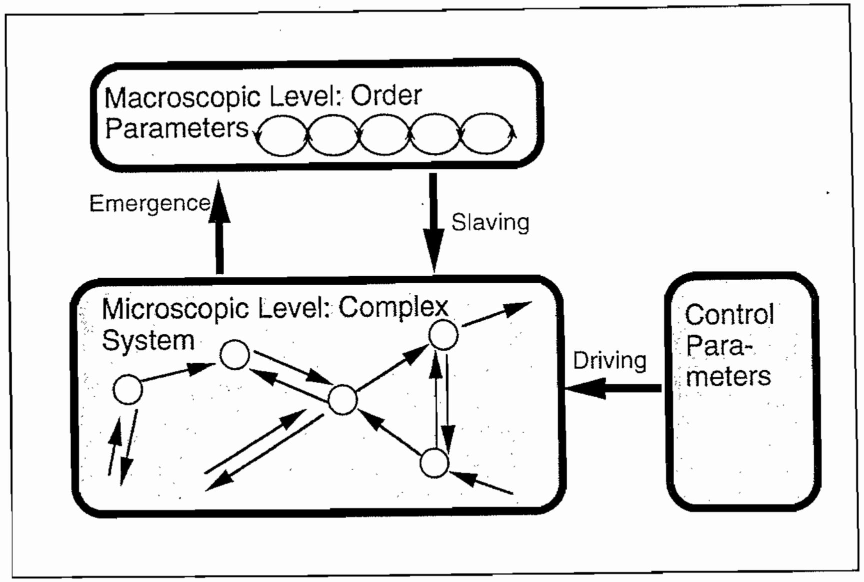

Figure 4: General model of a self-organizing system (schematic)

Application to cognition

Our suggestion is to apply what has been said about self-organizing systems in general to the cognitive-emotional system (CES) mentioned above. If the prerequisites for self-organization are met by the CES, this should yield an

outline of a self-organizational theory of cognition (Haken & Stadler, 1990; Tschacher, 1997).

At the beginning of this section we addressed the conceptualization of the microscopic level of a CES. The cognitive components may be seen as 'behavior kernels' (Tschacher, 1997), i.e. hypothetical cognitive micro-items not directly accessible to introspection and experiment. If we tentatively cross the mind-body language border, the micro components may be identified as the activation of neuronal cell assemblies which translate, for instance, into time-dependent EEG potentials over brain tissue (for a synergetic theory of brain dynamics see Haken, 1996). But we shall refrain from raising the eternal issue of the mind-body interface here.

Conscious cognition (thinking, memory, intention), at any rate, is to be located at the macroscopic level of our model; we view cognition as an order parameter of the CES, a pattern of the cognitive system. In New AI, cognition may be designed accordingly as attractors in the phase space of a robot. The SMC 2 model of situated categorization (Scheier & Lambrinos, 1996) consists of a number of neural networks which are connected to the input (a CCD camera) and to the effectors of the robot (arm-gripper system). It turns out to be feasible to operationalize the synchronization of these networks as being equivalent to macroscopic pattern formation of a CES. This synchronization can be assessed using principal-component analysis and related measures.

Consequently, cognition does not result from mere associations of single cognitive items, i.e. in terms of AI cognition is not propositional. Cognition emerges from a multitude of synchronized items whose activity is selected in the CES owing to the control parameters in its environment. We previously posited that control parameters select order parameters in such a way as to reduce the environmental nonequilibrium. Optimality and fitness are thereby increased.

It is important to note that our basic construct, the complex system CES, is free from any intentionality. There is no volition, wish, or motivational variable in the CES per se. Behavior kernels are just the potential elements of what — after synchronization via self-organization — may evolve into perception, thought, emotion, intention, plan, and consequently, action. The CES represents only the prerequisites to think and act. All intentionality results from the interaction of the CES with its environment, which we conceptualize as a driving of the CES by control parameters. The equivalence to New AI, where intelligence arises from situatedness, may gradually become apparent.

The application of synergetics to cognition is symbolized by Fig. 5.

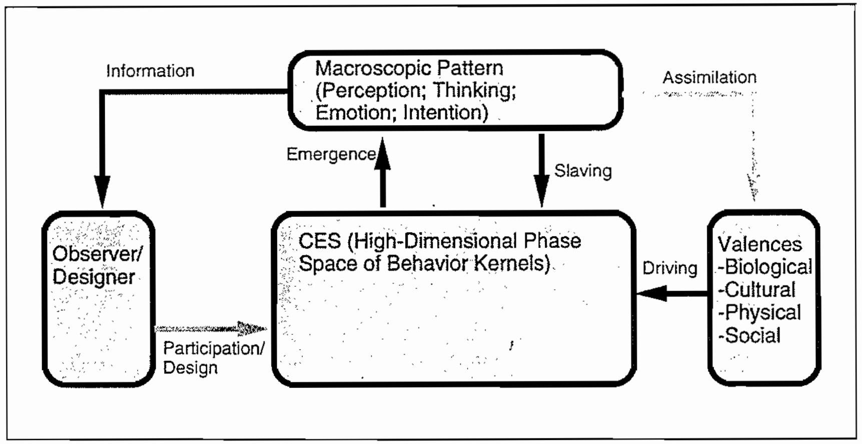

Figure 5: Schematic model of cognition as self-organized dynamics

Systems theory allows for a formulation of concepts suitable for a wide range of further applications. Particularly, a definition of 'complex system' can be tailored to the needs of *social theories*, as we may conceptualize variables of social interaction as the components of the system-under-study (Tschacher et al., 1992). One may then empirically test the occurrence of pattern formation and nonlinear phenomena against a different systems background, for example, social groups and therapeutic dyads (Tschacher & Brunner, 1995; Tschacher & Grawe, 1996).

Situatedness and motivation

The interaction of cognitive order parameters with the environment is the decisive point where synergetics turns into a situated theory; it will be shown how closely synergetic psychology and New AI are linked. We are first led to the question of what the concept of 'control parameter' refers to in the case of CESs and what 'nonequilibrium' means in cognition.

In psychology, energizing and incentive variables are treated under the heading of *motivation*. Several motivational theories have conceptualized motivation as reduction of libidinous tensions and disequilibria (as in Freudian psychology) or, in a more cognitive fashion, as attempts to reestablish cognitive balance and reduce dissonance (e.g. theories of social psychologists such as Heider (1958) and Festinger (1964)). In the latter's work, cognitive approaches go back again to Lewinian field theory (as does Gibson's concept of affordance mentioned in section 1).

Lewin's (1926) motivational construct of 'Aufforderungscharakter' or *valence* comes very close to our general idea of motivation being the driving of a CES by nonequilibrium control parameters. The difference lies mainly in the general theoretical frame: Field theory is based on the concept of forces ('vector psychology') as basic causes of psychological dynamics; valence then results from the tension in a psychological field. We opt for the synergetic approach, however, which is centered on dissipation and 'thermodynamic' nonequilibrium instead of forces. The progress achieved by this approach is that nonequilibrium dynamics (synergetics) can account for the formation of pattern and of attractors, whereas forces can only explain change. Thus, field theory must leave open the basic question of pattern formation and optimality (and for this reason cannot provide an answer to the symbol grounding problem). Nevertheless, the analogy to Lewinian thought is striking, so that we use and redefine the term 'valence' to denote a control parameter of a CES (see Fig. 5).

In Fig. 5 valences are located outside of the cognitive system CES. We recognize the sources of valence in the biological, social, cultural, and physical environment of the CES (seen this way, the body is 'environment' to cognition). The cognitive-emotional state of an individual is therefore continuously embedded in biological nonequilibrium (e.g. hypothalamic activity leading to 'hunger' cognition/emotion), cultural and social nonequilibrium (e.g. working atmosphere in an organization may facilitate or impede creativity), physical nonequilibrium (e.g. affordances built into housing and architecture, enabling certain 'standing patterns of behavior' while discouraging others (Barker, 1968)).

In reacting to the frame-of-reference problem, New AI argues for a wide definition of the cognitive system. Representation and memory, for example, also encompass the cultural setting and sensorimotor loops of an autonomous agent: culture and motor behavior can be seen as parts of cognition. Thus, there seems to be a discrepancy with our model which locates valences outside the CES. However this discrepancy is superficial and easily resolved: Valences are defined in such a way as to transfer pragmatic information (about various kinds of nonequilibrium) to the CES. A pebble is not a valence, but the pebble as an obstacle to an ant is (as well as its posing a challenge to a child!). Therefore valences are *interfaces* between pebbles and agents. We consider it to be of minor importance whether we view 'pebble-valence' inside cognition or as environment to cognition. We stress that a CES becomes *situated by valences*, and to us it seems purely a matter of terminology whether valences are seen as cognition or as constraints for cognition. Dynamical systems theory (e.g. Thompson & Stewart, 1993) is liberal when it comes to defining systems; what is conceptualized as part of, or environment to, a dynamical system-under-study is mainly a matter of convenience and convention. 'System' is not an ontological concept and should by no means be treated as such.

Evolutionary situatedness

New AI designs agents for specific ecological niches. In biological agents, it is obvious that the interaction of CES and valences is based on a long history of coevolution (i.e. joint and mutual evolution of CES and valences). For example, in human infants (but not in infant dolphins) there is a *preparedness* for visual cliffs, and for small, black, eight-legged animals: certain behavior kernels seem to exist from birth as predesigned candidates for being selected by certain valences in the physical environment.

The fit of system and control parameters / valences may take different forms:
(a) In the Bénard instability the fluid is chosen accordingly (gasoline would not work as well because of its viscosity and inflammability).
(b) In robotics, designers provide autonomous agents with predesigned value systems to adapt them for their niches.
(c) In the case of animals like the human individual, phylogeny has provided constraints for cognition by a highly prestructured (though not pre-wired!) neuronal substrate.

In the latter case (c) of 'natural' coevolution, the fit of system and control parameters is accounted for by the loop 'valence-CES-pattern-valence' of Fig. 5. This loop of coevolution deals with the symbol grounding problem; it has made the cognitive system capable of adaptation on a phylogenetic time scale. This time scale is much larger than the one of cognition in the here-and-now. In principle, however, the mechanism of coevolution is analogous to cognition seen as selection of order parameters in the here-and-now. Accordingly, Haken (1983) speaks of a "Darwinism of microscopic modes". Therefore, long-term coevolution is the platform on which self-organization occurs. Or rather, coevolution sets the stage for its own core method, namely self-organization (for a treatment of 'endosystems' — systems which modify their environment to modify themselves — see, Rössler, 1987; Atmanspacher & Dalenoort, 1994; Tschacher, 1997).

This is different in the two other examples given, (a) Bénard and (b) robot, where the observer or designer determines the systems' components and/or

valences in a meaningful way. The functioning and pattern formation in these systems depend on the participation of the observer or designer. We therefore have to introduce yet another loop (observer-CES-pattern-observer) into Fig. 5. This accounts for the frame-of-reference problem of New AI (Clancey, 1991), and more generally, for the observer-dependence of any observation. It seems that our descriptions can never be entirely free of an homunculus. Scientific explanation is always more or less threatened by infinite regress. The interactions related to this principle uncertainty are given by grey arrows in Fig. 5.

This point leads us to a caveat for robot design. The fundament of adaptivity in natural organisms is laid by coevolution. The adaptivity of cognition and action, with which this paper mainly deals, is grounded upon this fundament. Valences and cognition/action thus refer to each other because of their common history. If we design a robot to act autonomously in an environment, however, we shall have to design its value system (the valences) and its hardware substrate (e.g. the types and positions of sensors that cause a specific preparedness), at least to a certain extent. We should keep in mind, though, that the value systems and hardware should have the flexibility to evolve with the environment, because New AI knows well that pre-wired implementations lead to impasses as they confine the evolution of the system. Therefore, with pre-installed values and preparednesses an autonomous agent may have but few options to develop intelligent behavior; a competing computational expert system may consequently have an advantage right from the start (its designer did his or her best to create an optimal knowledge base). It is not easy for a synthetic approach like AI to beware of the homunculus of the designer.

The lesson of situated and self-organizing cognition at this point would be to provide an autonomous robot with many degrees of freedom, but just enough structure to get along and *organize itself*. We should heed Brooks' warning that it took evolution the longest to reach the simplest level of intelligence. The evolution of insect-level intelligence lasted 3 billion years, the subsequent evolution of human intelligence "only another 500 million years" (Brooks, 1991). A fascinating empirical question is how much faster AI might be.

The frame problem, the symbol grounding problem, and the frame-of-reference problem of AI are highly relevant for psychology in that they show that intelligence and adaptive cognition cannot be computed and implemented directly. The lesson of New AI for action psychology is that one-sidedness leads to pseudo explanations. New AI suggests that intelligence is perhaps unlikely, but may evolve from the interaction of several environments with and in a situated cognitive system. Since we know that intelligent cognition exists, we may suppose it is a result of evolutions at different time scales; it dwells at the interface of these evolutions.

References

Abraham, R.H.; Shaw, C.D., 1992, *Dynamics - The Geometry of Behavior,* Redwood City, Addison-Wesley.

Ach, N., 1910, *Über den Willensakt und das Temperament,* Leipzig, Quelle und Meyer.

Anderson, J.R., 1983, *The Architecture of Cognition,* Cambridge, Harvard University Press.

Atmanspacher, H.; Dalenoort, G.J. (eds), 1994, *Inside Versus Outside,* Berlin, Springer.

Barker, R.G., 1968, *Ecological Psychology,* Stanford, Stanford University Press.

Bestehorn, M.; Fantz, M.; Friedrich, R.; Haken, H., 1993, Hexagonal and Spiral Patterns of Thermal Convection, *Phys Lett A,* vol. 174, pp. 48-52.

Brooks, R.A., 1986, Intelligence without Representation, *Artificial Intelligence,* vol. 47, pp. 139-159.

Brooks, R.A., 1986, A Robust Layered Control System for a Mobile Robot, *IEEE Journal of Robotics and Automation,* vol. 2, pp. 14-23.

Carver, C.S.; Scheier, M.F., 1982, Control Theory: A Useful Conceptual Framework for Personality-Social, Clinical, and Health Psychology, *Psychological Bulletin,* vol. 92, pp. 111-135.

Clancey, W.J., 1991, The Frame-of-Reference Problem in the Design of Intelligent Machines, in van Lehn K. (ed.), *Architectures for Intelligence,* Hillsdale NJ, Erlbaum.

Festinger, L., 1964, *Conflict, Decision, and Dissonance,* Stanford, Stanford Universiy Press.

Fodor, J.A., 1975, *The Language of Thought,* New York, Crowell.

Gibson, J.J., 1979, *The Ecological Approach to Visual Perception,* Boston, Houghton Mifflin.

Greeno, J.G.; Moore, J.L., 1993, Situativity and Symbols: Response to Vera and Simon, *Cognitive Science,* vol. 17, pp. 49-59.

Haken, H., 1983, Synergetics — An Introduction, Berlin, Springer.

Haken, H.; Stadler, M. (eds.), 1990, *Synergetics of Cognition*, Berlin, Springer.

Harnad, S., 1990, The Symbol Grounding Problem, *Physica D,* vol. 42, pp. 335-346.

Heckhausen, H.; Gollwitzer, P.M.; Weinert, F.E. (eds.), 1987, *Jenseits des Rubikon: Der Wille in den Humanwissenschaften,* Berlin, Springer.

Heckhausen, H.; Kuhl, J., 1985, From Wishes to Action: The Dead Ends and Short Cuts on the Long Way to Action, in M. Frese & J. Sabini (eds.), *Goal Directed Behavior: The Concept of Action in Psychology,* Hillsdale, Erlbaum, pp. 134-160.

Heidegger, M., 1962, *Being and Time* (Original: *Sein und Zeit,* 1927), New York, Harper & Row.

Heider, F., 1958, *The Psychology of Interpersonal Relations,* New York, Wiley.

Helson, H., 1933, The Fundamental Propositions of Gestalt Psychology, *Psychological Review,* vol. 40, pp. 13-32.

Janlert, L.E., 1987, Modeling Change — the Frame Problem, in Z.W. Pylyshyn (ed.), *The Robot's Dilemma. The Frame Problem in Artificial Intelligence,* Norwood NJ, Ablex, pp. 1-40.

Kahneman, D.; Slovic, P.; Tversky, A. (eds.), 1982, *Judgement under Uncertainty: Heuristics and Biases,* Cambridge, Cambridge University Press.

Kolers, P.A.; Smythe, W.E., 1984, Symbol Manipulation: Alternatives to the Computational View of Mind, *Journal of Verbal Learning and Verbal Behavior,* vol. 23, pp. 289-314.

Kratky, K., 1992, Chaos and Disorder, in W. Tschacher, G. Schiepek & E.J. Brunner (eds.), *Self-Organization and Clinical Psychology,* Berlin, Springer, pp. 88-101.

Kugler, P.N., Turvey, M.T., 1987, *Information, Natural Law, and the Self-Assembly of Rhythmic Movement,* Hillsdale NJ, Erlbaum.

Kuhl, J., 1992, A Theory of Self-Regulation: Action versus State Orientation, Self-Discrimination, and Some Applications, *Applied Psychology,* vol. 41, pp. 97-129.

Langley, P., 1983, Exploring the Space of Cognitive Architectures, *Behavior Research Methods & Instrumentation,* vol. 15, pp. 289-299.

Law, L.-G., 1993, Symbolic Processing vs. Situated Action: A Dialectical Synthesis?, *Research report,* no. 24, Universität München: Lehrstuhl für Empirische Pädagogik und Pädagogische Psychologie.

Lewin, K., 1926, Vorsatz, Wille und Bedürfnis, *Psychologische Forschung*, vol.7, pp. 330-385.

Lorenz, E.N., 1963, Deterministic Non-Periodic Flow, *J Atmos Sci,* vol. 20, pp. 130-141.

McCarthy, J.; Hayes, P.J., 1969, Some Philosophical Problems from the Standpoint of Artificial Intelligence, in B. Meltzer & D. Michie (eds.), *Machine Intelligence,* vol. 4, pp. 463-502.

Miller, G.A.; Galanter, E.; Pribram, K.H., 1960, *Plans and the Structure of Behavior,* New York, Holt, Rinehart, and Winston.

Neisser, U., 1976, *Cognition and Reality (Principles and Implications of Cognitive Psychology)*, San Francisco, Freeman.

Neisser, U., 1982, *Memory Observed: Remembering in Natural Contexts*, San Francisco CA, W.H. Freeman.

Neisser, U.; Winograd, E. (eds.), 1988, *Remembering reconsidered: Ecological and traditional approaches to the study of memory*, Cambridge, Cambridge University Press.

Newell, A.; Simon, H.A., 1976, Computer Science as Empirical Inquiry: Symbols and Search, *Comm. of the ACM*, vol. 19, pp. 113-126.

Pfeifer, R.; Scheier, C. (submitted), *Introduction to "New Artificial Intelligence"*, Cambridge, MIT Press.

Prigogine, L.; Stengers, I., 1984, *Order out of Chaos: Man's New Dialogue with Nature*, New York, Bantam Books.

Pylyshyn, Z.W. (ed.), 1987, *The Robot's Dilemma. The Frame Problem in Artificial Intelligence*, Norwood NJ, Ablex.

Quillian, R., 1968, Semantic Memory, in Minsky, M. (ed.), *Semantic Information Processing*, Cambridge MA, MIT Press.

Rosch, E.; Lloyd, B.B. (eds.), 1978, *Cognition and Categorization*, Hillsdale NJ, Erlbaum.

Rössler, O.E., 1987, Endophysics, in J.L. Casti & A. Karlqvist (eds.), *Real Brains, Artificial Minds*, New York, Elsevier, pp. 25-46.

Scheier, C.; Lambrinos, D., 1996, Categorization in a Real-World Agent Using Haptic Exploratlon and Active Perception, in P. Maes, M. Mataric, J.-A. Meyer, J. Pollack & S. Wilson (eds.), *From Animals to Animats*, Cambridge, MIT Press, pp. 65-75.

Simon, H.A., 1969 (second ed.), *The Sciences of the Artificial*, Cambridge MA, MIT Press.

Steels, L.; Brooks, R. (eds.), 1995, *The Artificial Life Route to Artificial Intelligence: Building Embodied, Situated Agents*, Hillsdale, Erlbaum.

Suchman, L.A., 1987, *Plans and Situated Actions - The Problem of Human-Machine Communication*, Cambridge MA, Cambridge University Press.

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

Thompson, J.M.T.; Stewart, H.B., 1993, A Tutorial Glossary of Geometrical Dynamics, *International Journal of Bifurcation and Chaos*, vol. 3, pp. 223-239.

Toda, M., 1982, *Man, Robot, and Society*, The Hague, Nijhoff.

Tschacher, W. (1997), *Prozessgestalten — Theorie, Methodik und Empirische Studien zur Selbstorganisation in der Psychologie*, Göttingen, Hogrefe.

Tschacher, W.; Brunner, E.J., 1995, Empirische Studien zur Dynamik von Gruppen aus der Sicht der Selbstorganisationstheorie, *Zeitschrift für Sozialpsychologie*, vol. 26, pp. 78-91.

Tschacher, W.; Brunner, E.J.; Schiepek, G., 1992, Self-Organization in Social Groups, in W. Tschacher, G. Schiepek & E.J. Brunner (eds.), *Self-Organization in Clinical Psychology*, Berlin, Springer, pp. 341-366.

Tschacher, W.; Grawe, K., 1996, Selbstorganisation in Therapieprozessen — Die Hypothese und empirische Prüfung der "Reduktion von Freiheitsgraden" bei der Entstehung von Therapiesystemen, *Zeitschrift für Klinische Psychologie*, vol. 25, pp. 55-60.

Vallacher, R.R.; Kaufman, J., 1996, Dynamics of Action Identification. Volatility and Structure in the Mental Representation of Behavior, in P.M. Gollwitzer & J.A. Bargh (eds.), *The Psychology of Action*, New York, Guilford, pp. 260-282.

Wertheimer, M., 1912, Experimentelle Studien über das Sehen von Bewegungen, *Z. Psychol.*, vol. 61, pp. 165-292.

Wilson, E.O., 1975, *Sociobiology (abridged edition)*, Cambridge MA, Harvard University Press.