

Does Representation Need Reality?

Rethinking Epistemological Issues in the Light of Recent Developments and Concepts in Cognitive Science

Markus F. Peschl

Dept. for Philosophy of Science, Univ. of Vienna, Sensengasse 8/10, A-1090 Wien, Austria
Email: franz-markus.peschl@univie.ac.at

Alexander Riegler

CLEA, Free University Brussels, rue de la Strategie 33, B-1160 Brussels, Belgium
Email: ariegler@vub.ac.be

Abstract · This paper discusses the notion of *representation* and outlines the ideas and questions which led to the organization of this volume. We argue for a distinction between the classical view of referential representation, and the alternative concept of system-relative representation. The latter refers to situated cognitive processes whose dynamics are merely modulated by their environment rather than being instructed and determined by it.

Introduction

How is “reality out there” represented in our heads? Does cognition work with information from the “outside” world? Is our thinking exclusively determined by the world, is it exposed to perturbations from an apparently endless environment, or does cognition actively generate and construct the “world”?

To approach these questions, let us start with Plato’s well-known allegory of the cave (*The Republic. Book VII*). In this allegory, prisoners are bound in a cave in such a way that they cannot turn their heads or move around. They can only see a wall in front of them. The light of a distant fire behind them casts shadows on the cave wall of themselves and other people wandering around. The prisoners have been restricted to this perspective since birth. Therefore, their only perception of themselves and their world is through the moving shadows on the wall. Hence, the prisoners perceive and take these shadows to be the actual objects in the world rather than recognizing them as mere shadows of the “real” environment.

Obviously, when looking at the problem of representation we face a similar situation: we are the prisoners having only a mediated access to the “real” world. How can we distinguish the “true” reality from the fake? Wittgenstein approached this question in his *Tractatus Philosophicus*: “In order to tell whether a picture is true or false we must compare it with reality.” (2.223)

This makes the problem of reference to a “real outer world” clear. From such a perspective we have to *distinguish* between the world as the domain of our experience (actuality, *wirklichkeit*; *W*) and the world as the domain of things in themselves (*realität*; *R*):

- *Realität*—from the Latin “res” (thing)—connotes the ontologically given environment every realist makes reference to. By “ontology” we refer to the philosophical tradition of claiming the existence and recognizability of an independent world outside, the existence of things in themselves (*Dinge an sich* in the sense of Kant)—like the “real” people whose shadows are perceived by the prisoners.
- *Wirklichkeit* (actuality) on the other hand—stemming from the German verb “wirken”, meaning to have an effect on—can be used to designate the “constructed” world in our minds, made up by our experiences and (genetic) predispositions.

The epistemologically most trivial version of the relationship between *W* and *R* suggests: $W = R$. Here, one assumes that an organism perceives its environment directly and free from distortion, and the world of experience is directly related to an objective world (*radical realism*).

This is only a common sense relationship and seems to be epistemologically not very plausible. In the following we introduce more sophisticated concepts that have been developed¹.

Referential representation

For the *classical representational theory*, *wirklichkeit* is a function of *realität*, $W = f(R)$. In this view, our experiences are a distorted and non-proportional image of *realität*. This theory is unsatisfactory insofar as it does not tell us a lot about the nature of the distortions and their relationship to the observer. Furthermore, it is—from a neuroscientific as well as system-theoretic perspective—not very plausible that the representation W depends only on the *realität* R . Such a view, i.e., $W = f(R)$, might imply some kind of distortion.² The representational relationship, however, remains basically referential (which seems to be contradictory to our experiences).³

Context-dependent representation

Within a *dialectic–materialistic world view*, the construction of *wirklichkeit* is established by an interaction between an observer and the observed (environment/phenomenon): $W = f(R, O, C)$. O denotes the properties of the observer and C the cultural–historical background. Such a perspective takes into account the properties of the observer in the sense that her activity of construction plays an active role in the formation of the representation. However, it may be difficult to sufficiently define the properties of C and O .

Self-referential representation

The concept of *cognitive self-reference* describes perception and representation as perception of rela-

tions. Stimuli are mere peripheral energetic conditions (i.e., perturbations P) for a semantically closed and self-organizing cognitive system. The structure of the cognitive system determines which structural configurations of its surroundings are perturbations to the system, and which are not. The idea is that the cognitive system is in a dynamical equilibrium. This means that the perpetually acting components of the system (e.g., neurons) respond solely to the activity of other components. There is no other way of influencing their state without destroying them, much as a detector of radiowaves displays activity only in the presence of waves of a certain length. From the perspective of an observer, an organism (with its self-referential cognitive equipment) is embedded within its environment. Some of the processes within the environment (and external to the organism) act as perturbations to the organism: Through the sensory surface they are transduced into neuronal activity which in turn may have an impact on the dynamical equilibrium of the cognitive system. Generally speaking, cognitive self-reference means that *wirklichkeit* is a function of three variables: $W = f(W, E, P)$. E denotes the individual background experience of a specific organism; from the perspective of systems theory or computational neuroscience, E refers to the structure of the state space. In other words, E determines the space of possible successor states of every particular state of the system.⁴ Note that R is not part of the equation, as the neurons refer only to each other! In a constructivist context, we interpret the absence of R as a consequence of the *operational closure* of the cognitive apparatus. That is, the states of neural activity always originate from and lead to other states of neuronal activity in a recurrent, self-referring manner (Maturana & Varela 1979; Winograd & Flores 1986). From this point of view, to deal with an object means for an organism to deal with its own internal states.⁵ Later in this paper we will elaborate the idea of cognitive self-reference and

¹ A similar categorization can be found in Stadler and Kruse (1990).

² The distortion is caused by the function $f(\cdot)$ which can be thought of as a description of the distortion taking place, for instance, in the sensory system or in the processes occurring in the transduction or primary processing of the environmental signal.

³ Think, for instance of the phenomenon of color constancy; more generally speaking, think of all the situations in which *one* particular environmental event/phenomenon is experienced in *different* ways (according to our present internal state).

⁴ Of course, this is a *descriptive* perspective. Dynamically speaking, the nervous system does not refer to past states, e.g., its state 10 seconds ago (cf. the non-temporal characterization of the nervous system, Maturana & Varela 1979).

⁵ This of course does not refer to a solipsistic world view, since this equation describes a mapping onto *wirklichkeit* W rather than *realität*!

argue that it supersedes the prevailing referential concept of representation in the cognitive sciences.

At this point, an interesting question appears: What is the nature of the perturbations P in the cognitive self-reference perspective? This question gives way to many interpretations. A rather cautious position (e.g., Roth 1994) suggests that perturbations are certain patterns of energy transduced by the sensor organs that give rise to a mosaic of elementary events that the brain tries to make sense out of. This would mean that objects in the traditional sense (chairs, mountains, etc.) do not exist as objective entities but energy (such as electromagnetic waves) does (cf. Chandler's rejection of this concept, this volume). A somewhat more radical interpretation (such as Glasersfeld's) acknowledges the arbitrariness of accepting the absolute existence of energy while dismissing "ordinary" objects. Let's investigate this a little deeper. Above we emphasized that this view is the view of an observer who distinguishes between an organism and its environment, both being external to herself. However, we arrive at a crucial point when speaking of one's own cognitive self-reference: An observer who observes herself can never transcend this perspective. That is, we can place both ourselves and our environment only "external" to ourselves when describing their relationship. Therefore we cannot help but assume that there is indeed an external reality which is the originator of perturbations that influence our nervous system through mediation of our sensors. However, as we always have to do the somewhat bizarre step of thinking of ourselves as a third person, such a perspective can never reveal the "true" nature of perturbations, even worse, we cannot even claim nor deny the existence of reality.

Implications

What are the implications for our original question of whether representation needs reality? Obviously, the solution to the problem of appropriate representation depends on the perspective from which we look at the agent. One well-known example for a, cognitively speaking, "misleading" approach are microworld models used in artificial intelligence (Dreyfus & Dreyfus 1988): i.e., the attempt to determine basic cognitive mechanisms by reducing the complexity of the real world to the simplicity of a toy world and—starting from this simplified and "cognitively preprocessed world"—

to build up a representational system as a network of propositions. It is obvious that all the task of "understanding" was actually done by the programmers themselves. They put a large amount of pre-processed information into the system, thus creating a universe of interrelated facts (i.e., purely syntactical structures), but they did not create a meaningful world. As one can see in highly complex expert systems, such a strategy is not limited to small toy worlds, but can go very far on the level of the complexity of knowledge; it has to be clear, however, that this immense and quantitative increase in complexity (i.e., more rules, facts, relations, etc.) does not imply a new level (e.g., "semantics"). Hence, it does *not*—from an epistemological perspective—bring about a new quality.

Not only has symbolic cognitive science come to a dead end (in the context of the effort to explain cognitive processes). The euphoric days of the new possibilities of *connectionist networks*, including the simulation of cognitive activities and learning, have passed, too. It seems that cognitive science has returned back to—in Kuhn's sense (1970)—"normal science" or to "puzzle solving". As an implication of this development, connectionist cognitive science has almost lost itself in (technical) details, such as learning factors, minimal adjustments, and optimizations in learning algorithms or activation functions, and so on. In tackling only these technical "micro-problems", it seems to have failed to address the really interesting, pressing, qualitative, and "big" questions about cognition.

Nevertheless, in the context of recent developments in cognitive science (e.g., computational neuroscience, artificial life, autonomous agents, situated action, robotics, etc.) a considerable amount of epistemological potential can be found; basic issues have received *new inputs* and *new interpretations* over the last years. The goal of this volume is to investigate some of these trends and to make them more explicit in order to achieve some clarity as to where cognitive science might develop in the future. Among these issues are:

- The necessity of rethinking the concept of *representation* in the light of dynamical, top-down, and recurrent processes in natural and artificial cognitive systems, in particular in neural systems;
- The *embodiment of knowledge* in its substratum, e.g., in a neural structure;

- The importance of *simulation* as methodological tool for theory development in cognitive science and for a more profound conceptual understanding of cognition.

Among others, these issues have impacts on

- the evolution of semantics and symbol grounding;
- the design of autonomous systems acting in the “real” world, either as robots or software agents;
- the empirical research in (cognitive) neuroscience (as far as their assumptions, experiments, and interpretation of data goes);
- the philosophical interpretation of models being proposed by cognitive science;

In the following sections a discussion will be given on these points mentioned above. This is done in order to understand the “conceptual atmosphere” and the motivation for this volume.

Reviewing the Notion of Representation

The understanding of the concept of representation has received new stimuli from the developments in connectionism/computational neuroscience as well as in (empirical) neuroscience.⁶ The findings and concepts stemming from these fields seem to seriously question the traditional understanding of representation. There are at least three points which seem to be of great importance:

- (a) giving up linguistic transparency in favor of distributed representation;
- (b) giving up the classical referential view of representation in favor of a generative paradigm and the concept of functional fitness;
- (c) the concepts of embodiment and construction in knowledge representation.

Linguistic transparency

The classical understanding of representation is largely based on the idea that propositions (e.g., Fodor 1981) represent the (internal and external) environment in a more or less linguistically transpar-

⁶The terms connectionism and computational neuroscience will be used synonymously throughout this paper—for the purpose of the arguments to follow the difference between these two terms can be neglected.

ent way. With the advent of “truly distributed” connectionist models⁷, linguistic transparency had to be given up in favor of the concept of distributed representation and *subsymbolic* representation in microfeatures (cf. Singer’s contribution, as well as Gelder 1992, Hinton et al. 1986, Rumelhart et al. 1986, Smolensky 1988, and many others).

On a superficial level, giving up linguistic transparency seems to be one of the big disadvantages of distributed representation (see also Fodor & Pylyshin 1988; and many others), as it implies that we have to give up the traceability of the system’s dynamics as well. Taking a closer look reveals, however, that (a) it is by no means clear why our brain does its work by making use of the same semantic categories as our language does. (b) What is referred to as a proposition is the result of extremely complex processes occurring in the neural dynamics and leading to the externalization of “propositional categories” (e.g., in form of symbols, language, etc.). Hence, it is not at all clear why the processes responsible for generating these linguistic categories necessarily have to be based on and have to rely on exactly these categories. From the perspective of philosophy of science it seems rather questionable, if it is just to use in an explanation of so-called higher cognitive abilities (such as language) the same mechanisms/structures both in the explanatory mechanism/explanans (e.g., symbol processing mechanisms) and in the resulting behavior/explanandum (e.g., linguistic structures). (c) Furthermore, it is known from empirical neuroscience that the activity of most neurons cannot be explicitly related to semantically transparent phenomena or events.

All these considerations lead to the conclusion that the criterion of linguistically transparent representation is—perhaps—the result of our common sense experience, “auto-introspection”, and our common sense assumptions about representation. It seems that it arises from our need to somehow box everything and every process into linguistic categories (and, thus, also these processes which are leading to these categories) in order to make it available for cognitive manipulation. One of the objectives of this book is to investigate, if there is a necessity of linguistic categories for explaining cognitive phenomena and to search for alternative views.

⁷“Truly distributed representation” is mainly due to weight configurations which are the result of learning algorithms, such as the backpropagation learning rule.

Referential representation in recurrent architectures?

The second and epistemologically more important implication of connectionist systems concerns the very concept of representation: due to the highly *recurrent architecture* in the brain we are confronted with a phenomenon which questions the classical view of referential representation, i.e., $W = f(R)$: here, any state within the cognitive system refers to an (internal or external) environmental state in a more or less stable manner (independently whether it happens in a linguistically transparent manner or not). The recurrent architecture implies, however, that the neural system finds itself in a certain *internal state* at any point in time.⁸ This internal state has an indirect or direct influence on the following internal/representational states, as the resulting *top-down* processes are feeding back on the incoming stimuli.

This implies that the environmental state does *not* necessarily determine the representational (= internal) state any more, because the present internal state has to be seen as some kind of disposition for the possible successor-state. More precisely, the present internal state determines the space of possible representational successor states and the current environmental state only chooses one out of them. Hence, a different present internal/representational state determines a different space of possible successor-states which implies that the same environmental state might lead into different representational states (see also Peschl 1997). In other words, a single environmental state/event can be represented in different representational states (depending on the precedent internal state). This implies that the concept of a stable referential relationship of representation can *not* be found any more.

In other words, the classical idea of an environmental state determining the internal/representational state has to be forsaken due to the feedback influence of previous internal/representational states in the recurrent neural system. The idea of a more or less stable relationship between a representational state and an (internal or external) environmental state has to be abandoned (for further details see Peschl 1997). Rather, the influence of the environ-

⁸ This also applies to feed forward architectures; however, in these architectures the internal state does not have an influence on the incoming activations because the internal state is “shifted out” of the network in each time step.

mental input has to be reduced to the *modulation* of the internal representational dynamics (*perturbations* P as discussed above). Unfortunately, the importance of this far reaching epistemological issue has not been acknowledged by a large group within the cognitive science community.

Representation, construction, and generation of behavior

So, if the aim of representation is no longer to map the environment as accurately as possible we have to characterize it as the *generation of behavior* in terms of *functional fitness* (cf. contribution of von Glasersfeld who speaks of the *viability* of representations). I.e., behavior which (i) facilitates the organism’s survival (in the broadest sense) and (ii) functionally fits into the particular environmental context.⁹ It is therefore no longer necessary to search for neurons (or groups of neurons) whose activations correlate with external events in a stable referential manner.

As the representational structure is the result of a system-relative construction process, it is no wonder that we are experiencing difficulties identifying (traditional, referential) representations in natural and artificial neural systems. It seems that it is simply the wrong thing to search for. Understanding representation from the perspective of constructivism and the concept of functional fitness (e.g., Glasersfeld 1984, 1995) gives us a clue as to what we have to look for in the representational substratum; namely, mechanisms which allow the generation of adequate behavior. Traditionally these mechanisms have been thought of in terms of manipulations on referential representations, however, there is neither neuroscientific nor epistemological evidence in favor of such a view.

In this context it is important to note that a similar problem arises in most approaches in artificial life. As argued in Riegler (1997), a typical deficiency of many artificial life models is the *PacMan syndrome*: Simulated organisms interact with anthropomorphically defined entities, such as “food” and “enemy”. Such models perform a mere optimizing task which yields a maximum gain of energy together with a minimum loss of health. No attention is paid to

⁹ Recent developments in agent-based architectures speak of situatedness of reactive agents (cf. Clancey 1997).

questions like: How have organisms arrived at the idea that something is a source of food? How do they “know” that another creature is a dangerous opponent? Predators do not run around with a label saying “I’m your enemy”. Even if this would be the case—how would have cognitive beings learned to understand the meaning of those labels? (Think of the difficulties to understand signs in a country whose language and letters you don’t know at all.) However, things seem to be different within mathematical models. If we look at mathematical formulae we (usually) know what the meaning of the labels (variables) is although we are not the author of the equations. E.g., arriving at a result of $m = 8.3$ we know what m is, how meaning got attached (namely a-priori to the calculation). In logical calculus, semantics defines meaning and truth in terms of an underlying model, ontology, or logical interpretation. Cognitive sciences, however, transcend this purely symbolic framework. While we—as designer of artificial life models—would like the upper left pixel on the computer screen to be a food pill for the pixel in the lower right corner representing the cognitive creature, this is not necessarily true from the perspective of the creature and its cognitive apparatus. The question regarding the phylogenetic and ontogenetic emergence of a system-relative representation, i.e., “meanings” for the organisms, is not touched.

Furthermore, it is important to note that the dynamics of a recurrent cognitive architecture need not necessarily be implemented in a typical neural network manner. Riegler (1994, 1997) describes the implementation of a rule-based system: the algorithm operates exclusively on a set of “internal state cells” (having no explicit linguistic reference) rather than on anthropomorphically predefined sensor- and motor-states. As this preserves the idea of operational closure (Maturana and Varela 1979) it is, too, an implementation of cognitive self-reference and hence transcends purely a referential representation. Thus, we conclude that there is no epistemological difference between rule-based systems and connectionist approaches as long as we avoid forcing a referential representation scheme.

Embodiment, construction, and dynamics of knowledge

In the context of these questions one can observe a shift in interests in the field of knowledge represen-

tation in the last decade: the focus has changed from trying to capture or depict environmental structures statically onto the representational structure towards stressing more the question of the genesis, *development*, and *dynamics* of knowledge. Connectionist approaches (i.e., their focus on learning strategies), genetic algorithms, and the combination of both (e.g., Elman et al. 1996, Cangelosi et al. 1994, and many others) had a crucial impact on the development of representation mechanisms modeling learning and the dynamical aspect of knowledge. Furthermore, results from empirical neuroscience (sometimes having been triggered by computational approaches and concepts) have brought about a better understanding of the learning mechanisms which are responsible for the dynamics of knowledge in our brains¹⁰.

One of the most important epistemological implications of this development and the above discussions is the insight that *knowledge is the result of an active construction process* rather than of a more or less passive mapping (cf. Sjölander, this volume). In this view, the organism actively extracts and constructs these environmental regularities which are relevant for its particular survival. The environment no longer instructs or determines the structure of the representation system, but only plays the role of *constraining* the construction processes. In other words, the knowledge can be freely constructed as long as it does not “violate” the environmental constraints. The result is a *system-relative* representation of knowledge (about environmental regularities), where $W = f(W, E, P)$, as presented in the introduction: representation does not need reality as an instructive instance! Hitting a fly with a flap yields a different effect than beating an elephant with the same flap: Representation depends on the structure of the cognitive system rather than on outside entities.

In this context it seems that we need to take the concept of *embodiment of knowledge* more seriously. Knowledge (representation) can—at least in neurally based cognitive systems—no longer be understood as something abstract and completely

¹⁰ Unfortunately, we are standing only at the very beginning to fully understand these processes. However, the basic principles (e.g., long term potentiation/depression (LTP, LTD), Hebb’s concepts, etc.) seem to be quite promising and have brought about a new understanding of knowledge.

detached from the (neural) substratum. Rather, we have to make the effort and try to understand neural structures, architectures, and dynamics in terms of contributing to the production of functionally fitting behavior on a non-referential basis of representation. In other words, the so-called environmental regularities are not stored explicitly in the structure of the synaptic weights. A particular neural architecture has to be understood as the result of a long phylo- and ontogenetic adaptation/construction process which aims at relating the organism-relevant environmental regularities with the organism's requirements for production of behavior ensuring its survival.

At the heart of these construction processes, neural learning, adaptation, and plasticity, as well as phylogenetic processes can be found. All these processes are the substratum for any learning dynamics continuously occurring in a cognitive system. Consequently, knowledge (representation) in a cognitive system cannot be considered something static, but has to be seen as a highly dynamical process continuously adapting to the changing (internal and external) environmental constraints. One of the intents of this volume is to study exactly this relationship between the neurophysiological processes (of learning and adaptation) and its epistemological implications.

Simulation as Methodological Tool

The last years have shown a sharp increase in the importance of the method of simulation in the context of theory development in cognitive science. The extensive use of simulation brings about a whole new methodological approach and dynamics in disciplines which formerly were working almost exclusively empirically, such as neuroscience, psychology, biology, and physics. The interesting insights which are achieved by simulation are not so much results about details, but concern *conceptual* knowledge which can be used as input and *stimulation* for both empirical and epistemological investigations.

One of the main purposes of psychology, (cognitive) neuroscience, linguistics, and many other "cognitive disciplines" has always been a better understanding of so-called *cognitive* processes. Most of the resulting approaches to cognition were based on empirical investigations and/or more or less speculative and common-sense interpretations

of cognitive phenomena. Progress in empirical sciences is based on a continuous process of construction, negotiation, and adaptation to the "empirical data". The target of this process is to reach a state of (epistemological) equilibrium in which the theory fits into the environmental dynamics, meaning that the theory—at least—predicts the environmental dynamics correctly within some margin of error. Often the complexity of cognitive processes and their substratum does not match the comparably poor empirical approaches and understanding of cognitive phenomena (cf. Dorffner's contribution which stresses the importance of connectionism as a helpful modeling framework to understand cognition). Therefore, much room is opened up for rather speculative concepts in this field.

Fortunately, the simulation method introduces a new dimension to cognitive science and, more specifically, to computational neuroscience/connectionism. Simulation models are especially interesting in the context of cognitive neuroscience, as its empirical results and theories are sometimes so rich in detail (e.g., data on the release of neurotransmitter, theories on a molecular level, etc.) that it is almost impossible to relate them to cognitive phenomena. In other words, there is an explanatory *gap* and a strong tension between the epistemologically inspired questions on cognition (e.g., about knowledge representation) and the empirical and highly detailed results from neuroscience. In this context the connectionist approach—in the broadest sense—plays a crucial role as *mediator*: it stands between the two poles of the rather speculative epistemological theories and the empirically grounded neuroscientific details and—in many cases—makes them compatible. This *compatibility* is achieved by the trick of focusing on the *conceptual level* of neural processes. By doing so, the most important characteristics and structures of neural systems, such as parallel processing, network architecture and massive connectivity, and distributed representation, are captured in a more or less simplified computational model whose dynamics can be related to and is directly relevant for epistemological and "cognitive" issues.

So, why do we stress the importance of simulation models of cognition in this book? It is not so much the technical details of simulation which we are interested in, but rather in the conceptual implications which these models have on the problem of knowledge representation. Hence, one of the objec-

tives of this volume is to show, how focussing on exactly this conceptual level can bring about *both* an empirically and epistemologically sound understanding of the ancient problem of representation in cognitive systems. Furthermore, including simulation techniques as a necessary tool for theory construction can guide empirical research not only on the level of technical details, but—and this seems to be even more important—on a conceptual level (e.g., concerning the assumptions/premises of a research strategy, the epistemological framework and foundations, etc.).

Conclusion

We have argued in favor of completely rethinking basic issues in cognitive science in the context of recent developments in this field. The main issue seems to be the question of knowledge representation, which changes dramatically with the advent of the concepts of connectionism and artificial life. We have suggested that the concept of a referential understanding of representation should be replaced by a system-relative form of representation which is not necessarily semantically transparent.

We argued from an epistemological and neuroscientific perspective that the task of *generating behavior* is more important than the accurate mapping of environmental structures to representational structures. It is by no means clear what the “point of reference” could be for an “accurate mapping”. Is it our own perception and conceptualization of the world, or that of a rat, or the world itself,...? One is tempted to assume that the outside world (in the sense *realität*) acts as some kind of constraint for our construction/representation processes. But assuming the existence of that *realität* (or parts of it such as electromagnetic waves; see above) would put us into the camp of referential realists as we (though not trivially) map our experience onto the idea of an outside world. Assuming the non-existence of the world, on the contrary, would stigmatize us as solipsists. As a solution we have therefore to accept that claims about the existence of an objective *realität* might not be necessary for the purpose of scientific explanations.

We have seen that such a view of knowledge representation is closely related to constructivist concepts. In this approach to epistemology, the only criterion for successful knowledge (representation) is its functional fitness. Furthermore, it has become

clear that knowledge is not a static structure, but is continuously changing. This dynamics can be described as a process of construction and adaptation and finds its substrate in the neural dynamics/plasticity.

The main goal of this volume is to discuss these fundamental shifts in cognitive science and to sketch the implications on an epistemological and methodological level for cognitive science and its related disciplines.

References

- Cangelosi, A., Parisi, D. & Nolfi, S. (1994) Cell division and migration in a genotype for neural networks. *Network: computation in neural systems* 5(4): 497–516.
- Clancey, W. J. (1997) *Situated Cognition. On Human Knowledge and Computer Representations*. New York: Cambridge University Press.
- Dreyfus, H. L. & Dreyfus, S. E. (1988) Making a Mind versus Modelling the Brain: Artificial Intelligence Back at a Branch-Point. *Artificial Intelligence* 117: 309–33.
- Elman, J. et al. (1996) *Rethinking innateness. A connectionist perspective on development*. Cambridge, MA: MIT Press.
- Fodor, J. A. (1981) *Representations: philosophical essays on the foundations of cognitive science*. Cambridge, MA: MIT Press.
- Fodor, J. A. & Pylyshin, Z. W. (1988) Connectionism and cognitive architecture: a critical analysis. *Cognition* 20. Reprinted in: Beakley, B. et al. (eds.) (1992) *The philosophy of mind*. Cambridge: MIT Press.
- Gelder, T. von (1992) Defining “distributed representation”. *Connection Science* 4(3/4): 175–191.
- Glaserfeld, E. von (1984) An introduction to radical constructivism. In: Watzlawick, P. (ed.) *The invented reality*. New York: Norton, pp. 17–40.
- Glaserfeld, E. von (1995) *Radical constructivism: a way of knowing and learning*. London: Falmer Press.
- Hinton, G. E., McClelland, J. L. & Rumelhart, D. E. (1986) Distributed representations. In: Rumelhart, D.E. & McClelland, J. L. (eds.) *Parallel Distributed Processing: explorations in the microstructure of cognition. Foundations*, Volume I. Cambridge, MA: MIT Press, pp. 77–109.
- Kuhn, T. S. (1970) *The Structure of Scientific Revolutions* (2nd ed.) Chicago: Univ. of Chicago Press.

- Maturana, H. R. & Varela, F. J. (1979) *Autopoiesis and Cognition: The Realization of the Living*. Boston: Reidel.
- Peschl, M. F. (1997) The Representational Relation Between Environmental Structures and Neural Systems: Autonomy and Environmental Dependency in Neural Knowledge Representation. *Non-linear Dynamics, Psychology, and Life Sciences* 1(2): 99–121.
- Riegler, A. (1994) Constructivist Artificial Life: The constructivist–anticipatory principle and functional coupling. In: Hopf, J. (ed.) Proceedings of the 18th German Conference on Artificial Intelligence (KI-94) Workshop on Genetic Algorithms within the Framework of Evolutionary Computation. Max-Planck-Institute Report No. MPI-I-94-241, pp. 73–83.
- Riegler, A. (1997) Ein kybernetisch–konstruktivistisches Modell der Kognition. In: Müller, A., Müller, K. H. & Stadler, F. (eds.) *Konstruktivismus und Kognitionswissenschaft. Kulturelle Wurzeln und Ergebnisse*. Wien, New York: Springer, pp. 75–88.
- Roth, G. (1994) *Das Gehirn und seine Wirklichkeit*. Frankfurt: Suhrkamp.
- Rumelhart, D. E., P. Smolensky, McClelland, J. L. & Hinton, G. E. (1986) Schemata and sequential thought processes in PDP models. In: McClelland, J. L. & Rumelhart, D. E. (eds.) *Parallel Distributed Processing: explorations in the microstructure of cognition. Psychological and biological models*, Volume II. Cambridge, MA: MIT Press, pp. 7–57.
- Smolensky, P. (1988) On the proper treatment of connectionism. *Behavioral and Brain Sciences* 11: 1–74.
- Stadler, M. & Kruse, P. (1990) Über Wirklichkeitskriterien. In: Riegler, V. & Vetter, C. (eds.) *Zur Biologie der Kognition*. Frankfurt a. M.: Suhrkamp, pp. 133–158.
- Winograd, T. & Flores, F. (1986) *Understanding Computers and Cognition: A New Foundation for Design*. Norwood, NJ: Ablex.

