Riegler, A. (2005) Cognitive systems do not pick up knowledge from the environment. In: Russell, A. et al. (eds.) Proceedings of AMKLC'05, International Symposium on Adaptive Models of Knowledge, Language and Cognition Helsinki University of Technology, Espoo, Finland, pp. 1–7.

COGNITIVE SYSTEMS DO NOT PICK UP KNOWLEDGE FROM THE ENVIRONMENT

Alexander Riegler

Center Leo Apostel for Interdisciplinary Research, Vrije Universiteit Brussel
Krijgskundestraat 33, B-1160 Brussels, Belgium, ariegler@vub.ac.be

ABSTRACT

The conventional philosophical perspective on knowledge and meaning suffers from at least two problems: the proper definition of truth, and the use of propositions as a basic vehicle for knowledge. By contrast, naturalized accounts such as evolutionary epistemology and radical constructivism offer a broader conceptual frame which allows to include the cognition of animals and artifacts as well. The paper explores a possible mechanism, the cognitive psychologist notion of schemata-controlled information pickup against a radical constructivist backdrop. I outline the importance of embodiment and autopoietic systems, which control their input rather than their output. The latter is considered the perspective of the observer-designer of cognitive artifacts. Using insights from philosophy and empirical results, the paper presents the implications of a radical constructivist understanding of knowledge and meaning. These include the rejection of premises evolutionary epistemology is based upon, and point in the direction of how to implement knowledge systems. The paper concludes with a call for closed-loop systems.

1. INTRODUCTION

In 1941, Konrad Lorenz believed to have solved one of the greatest philosophical challenges, viz. Kant's a prioris of space and time, which are considered fundamental to the faculty of knowing. Lorenz's solution to Kant's problem was straightforward: "Our categories and forms of perception, fixed prior to individual experience, are adapted to the external world for exactly the same reasons as the hoof of the horse is already adapted to the ground of the steppe before the horse is born and the fin of the fish is adapted to the water before the fish hatches" ([20], pp. 124–125). His solution is remarkable not only because it laid the foundation for what later became known as evolutionary epistemology (EE), but also for anticipating W. V. O. Quine's [34] demand to naturalize epistemology a few decades later.

The goal of my paper is to explore Lorenz's attempt to define knowledge as adaptation to the environment, in the light of another epistemology, radical constructivism (RC). Coincidentally, both epistemologies appeared on the scientific stage in the same year, 1974, and both seem to have many aspects in common.[1] Most prominently, they stress the role of the cognitive organism that actively acquires or constructs knowledge, respectively. Their naturalized account for knowledge and meaning is in contrast to more conventional philosophical arguments, which consider knowledge to be justified true beliefs (JTB). This traditional perspective suffers from the problem that its three components – justification, truth, and belief – do not seem to be sufficient to account for knowledge, as demonstrated by Edmund Gettier [11] and others who extended his original counter examples. From the angle of RC and EE, the difficulties of the conventional definition of knowledge can be traced back to at least two more elementary problems: (1) The lack of a proper definition of "truth" (i.e., when is a given proposition true?); (2) The restriction of the notion of knowledge to human cognition based on the assumption that knowledge must be formulated in the form of propositions.

The line of argumentation I chose to follow in this paper starts with a reflection on the notions of meaning and knowledge. I shall introduce cognitive concepts such as schemata-guided information pickup and ethological results pertaining to animal behavior along with philosophical thought experiments in order to arrive at a broader naturalized definition of knowledge. Furthermore, I will show how such a conception of knowledge and meaning is situated within the radical constructivist epistemology, and reject the perspective of EE. The paper concludes with implications for philosophy and the engineering of cognitive artifacts.

2. THE NOTION OF KNOWLEDGE PICKUP

In his 1985 book *Surely, You Are Joking Mr. Feynman!* [7], Richard Feynman warned the reader of the attempts by charlatans and astute businessmen to sell their ideas by giving them a pseudo-scientific appearance. He introduced the notion of "cargo cult science" on the basis of the story of the inhabitants on a fictive South Sea island. During World War II airplanes supplied them with goods. Consequently, the islanders wanted to have this

[1] Donald Campbell's 1974 paper "Evolutionary epistemology" [4] mentioned the notion the first time, and Ernst von Glasersfeld [12] did the same for RC in the same year.

happen again. So they started to create runways with fires along their sides; they set up a wooden hut for a man to sit in, with two wooden plates on his head as headphones, and bars of bamboo sticks looking like antennas. The form was perfect; everything looked the way it had been looking before. But, not surprising to us, it didn't work; no plane ever landed.

While Feynman's intention was to rant against pseudo science, he (accidentally or not) also covered the topic of the present paper. This becomes evident when we start posing questions such as: What did the islanders not understand? How could the meaning of food supply escape their understanding? And what was missing in their knowledge that made them build a superficial, appearance-oriented copy of an airport? We are judging these questions from the privileged position of being part of the civilization that came up with the idea of food supplies and airplanes in the first place. It comes at little surprise for us to learn about the existence of such concepts and entities because to us the story makes sense in the light of our past experiences. Embodied as we are in Western culture, we are prepared to assimilate the story into our existing network of experience, which includes encounters with planes and support for the Third World. It is an expression of the fundamental working of our cognitive system: A cognitive "schema accepts information as it becomes available at sensory surfaces and is changed by that information; it directs movements and exploratory activities that make more information available, by which it is further modified" ([28], p. 55). This amounts to a mutual interplay between cognitive apparatus and the information it retrieves. Consequently, "information" makes only sense to the individual who can integrate it into the existing network of schemata. The network itself may undergo modifications due to the integration of the new experience. Both assimilation and modification (accommodation) of cognitive structures were already postulated by Jean Piaget in the first half of the 20th century. On a more abstract, philosophical level, a similar thought was also expressed by Italian Renaissance philosopher Giambattista Vico. His assertion "verum ipsum factum"[2] expresses the idea that, prior to any act of knowing, conceptual structures must be erected that can pick up new knowledge. In the rest of the paper, I will detail how such schemata-controlled "information pickup" addresses the problem of knowledge.

3. MEANING IN NATURAL AND ARTIFICIAL SYSTEMS

Let us investigate Vico's claim in a more technological context. His original intention was to claim that God understands nature, while man understands man-made structures. By logical entailment, this provokes the following challenging questions: When you turn on your computer, do you as a computer user understand the working of that device? If you are a programmer and make the machine carry out a certain sequence of computational steps according to your will—like the islanders wanted the planes to come back and to land on their island—do you understand the device? If you are an electronic engineer and about to design a new integrated circuit which forces electric currents to follow certain tracks and to activate certain electronic components, do you understand the machinery you created?

These questions suggest that no matter at which level you are, you can never *fully understand* the meaning of the phenomenon you are faced with although you may feel *fully satisfied* by your explanations. Piaget [31] refers to it as "assimilation… to the desires and affectivity of the subject" and Frank Keil [17] calls it the "sense of comprehension" we are in need of. His empirical findings reveal that people greatly overestimate their knowledge of facts and procedure. This illusion of explanatory depth, however, is little better than ignorance. Proponents of the phlogiston theory, for example, were not a bunch of cranks. They seriously tried to give coherence and sense to their observations and created the first comprehensive chemistry theory. But as careful experimentation revealed, they did not come to know the 'accurate' explanatory mechanisms to account for combustion. They happened to try one of the virtually infinite possible mappings from the appearances to the inner working of natural entities but they were as wrong as a randomly guessing amateur. So is knowledge nothing but an accidental correspondence between the structure of the cognitive apparatus and the phenomenon in question?

To start with, knowledge does not necessarily refer to understanding directly observable processes. For cognitive systems it is quite important to get a grip on mediated forms of knowledge such as linguistic utterances, both in written and oral form. The following thought experiment demonstrates that this is no trivial matter. In the Western hemisphere, most people certainly know what a "mermaid" is even though chances are small that they have ever met one. Let us assume for a moment that it is the first time you hear the word "mermaid." You are told that it is a hybrid creature between woman and fish. Therefore, you are at ease to construct a representation out of already known elements that are associated with "woman" and "fish," i.e., a composite which is a fish-tailed biped. Of course, as such it does not resemble the imaginary creature of the sea. Your deviant notion might never get into troubles with the stories you read afterwards. The plot of those stories is compatible with both variants unless you encounter a picture of a mermaid. With this thought experiment, Ernst von Glasersfeld [13] intended to show that one will modify the concept that is the subjective interpretation of the word (or, in terms of more conventional philosophy, one's belief) only if some context forces one to do so. However, the relation between cognitive structure and phenomenon seems to be arbitrary; as long as the former remains viable with re-

[2] "The truth is the same as the made" in the sense of "we know only what we put together."

gard to the action of its owner it does not need to bear any structural resemblance with what it is supposed to represent. It is sufficient that some cognitive schemata pick up the right chunks of information. Such a situation can be compared with the working of a master key the notches of which provide the lock with all the necessary positions the lock's pins ask for. Insights from ethology indeed teach us that animal behavior is triggered by such "master key" stimuli. This can be shown in experiments using dummies, i.e., copies of natural key stimuli reduced to a few crucial features. For example, young blackbirds open their beaks if simple dummies are presented. Juvenile gulls mistake a pointed stick for an adult bird. Dummies that do not even closely resemble the appearance of the animal cause aggressive behavior in male sticklebacks. In all these cases, the dummies, although they do not resemble the "real thing," have a meaning for the animal. From the perspective of conventional philosophy, however, the young birds do not know anything although they may believe that they are approached by the parenting bird, and this belief is certainly justified by the fact that hungry young birds usually get food. But then, given the evolutionary success of their behavior, what else if not knowledge has been transmitted from generation to generation? Despite the animals' susceptibility to rather inaccurate stimuli their knowledge is successful on an evolutionary scale. Furthermore, it is evident that correspondence with the real environment is not required for the generation of meaning. Rather, these examples suggest that meaning is an involuntary process that results from the interaction among systems.

4. THE LACK OF MEANING IN SYNTACTICAL SIMULATIONS

The question of how meaning is evoked in humans needs further clarification. Let us explore this question in more detail by comparing (1) the meaning of a complex mathematical formula such as a differential equation describing a continuous time dynamic process, and (2) the meaning of the dynamical behavior of pixels on a screen in a computer simulation such as the artificial life models in [5] or [6]. There, some pixels are supposed to represent ants (or people), which are engaged in some collective work. These pixels neither include any of the ants' physiological details nor do they resemble their appearance. Such simulations are deliberately designed by the programmer. They do not have the autonomy of natural systems. Therefore, it can be argued that this sort of simulation is just a computer game rather than a scientific representation of natural phenomena. The simulated creatures behave according to a priori specified rules rather than according to behavioral patterns resulting from phylogenetic and ontogenetic processes, as is the case with human beings, who, as Maturana [23] put it, "are the arising present of an evolutionary history in which our ancestors and the medium in which they lived have changed together congruently around the conservation of a manner of living in language, self-consciousness and a family life." In a certain sense, the behavior of simulated creatures is triggered by syntactical appearances rather than by an intrinsic link between creature and environment. As a result, natural ants, unlike their artificial counterparts, seem to understand their environment at the collective level rather than infer their behavioral repertoire from rules described in biological textbooks – as programmers might intend to do.[3] The crucial difference between natural ant and artificial ant is that the latter is designed according to the observable behavior of the former. This perceivable output of the natural ant, however, is the product of system-relative activities of the ant's physiology and (albeit miniscule) cognitive apparatus. The natural creature seeks to control its input (through perception and proprioception) rather than its output (e.g., [32]). Actions are just the means to control the input. They execute certain actions in order to change their input state, such as avoiding the perception of an obstacle or drinking to quench the proprioceptive feeling of thirst. Since the observer-designer cannot know what it is like to be an ant (its "first-person experience"), all she can do is to model the ant according to its observable behavioral pattern. She defines systems over the range of their behaviors and builds them accordingly, which results in anthropomorphic rather than autonomous artifacts. The islanders in Feynman's story are exposed to the same situation. They, too, model the visible state of affairs because they have no insight into the inner workings of airplanes and the socio-political context of the food supply—they are not *embodied* [38]. From the perspective of embodiment, the lack of understanding results from a lack of being embodied in the world of Western politics and technology. Is not this like mistaking the artifact (such as computer pixels) for the natural thing? Is not a simulation just as bad as the functionally worthless wooden equipment in the above example?

This bears a resemblance to what Jack Cowan (quoted in [16], p. 74) referred to as the *reminiscence problem* from which creators of computer simulations seem to suffer: "They say, 'Look, isn't this reminiscent of a biological or physical phenomenon!' They jump in right away as if it's a decent model for the phenomenon, and usually of course it's just got some accidental features that make it look like something." This syndrome is similar to the old philosophical conundrum of how to know that a model of a natural system and the system itself bear any relation to each other. How can a deductively working system, such as mathematics, allow for building bridges and flying to the moon?

The answer is that between understanding a mathematical formula and pixels on a computer screen, there is no difference. Rather, it is a matter of convention. Illiterates might find the visual presentation of computer simu-

[3] Also John Searle's Chinese room thought experiment is situated at this level of syntactical rule-following.

lations more attractive than the description in mathematical terms. Anybody familiar with mathematical notation will find it easy to understand the working of equations by means of an understanding of the relationships expressed by the mathematical operators.[4] A computer simulation does nothing else than establish a similar type of deterministic relationships among computational entities—just by using a different set of conventionally agreed symbols and instructions. Yet both mathematical equation and simulation amount to the same, i.e., a description of the output of the system. Neither of them, however, changes the fact that these systems are *not* embodied.

5. NATURALIZED EPISTEMOLOGIES AND THE INPUT–OUTPUT DICHOTOMY

At this point it is worth pointing out a major difference between radical constructivism and evolutionary epistemology. Being ethologists and morphologists, leading proponents of EE have focused in their publication on observable behavior (e.g., [21]) and biological forms (e.g., Riedl [35] referring to a "morphology of knowledge and explanation"), which gives rise to the assumption that they are interested in the *outside* view. They observe behavior and postulate a link between their rule-like behavior and general laws of cognition and knowledge acquisition. So their focus of attention is the *output* of the observed system, which they map onto their own experiential network. Referring to what has been said so far in this paper, behaviors are anthropomorphically[5] *attributed* in the following sense. As already pointed out for the context of artificial systems, an observer is not necessarily embodied in the world of the observed animal [27] and consequently interprets its behavior within her own referential system of understanding. As Pierre Duhem's concept regarding the underdeterminism of theories claims, even if we had the intellectual capacity for making inferences from the appearance to the inner working, we would face an innumerable amount of possible mappings from observational data onto the model because data points can be explained in any arbitrary way [26].[6] Facing this intellectual problem, all we can do is *trivialize* complex systems [9]. That is, we reduce the degrees of freedom of a given complex entity to behave like a trivial machine, i.e., an automaton without internal states.

In contrast to EE, proponents of RC have developed a preference for turning their attention to the mechanisms *inside* systems, i.e., their theories start with the *inner* perspective. A prominent example is Maturana's theory of autopoietic systems[7], the crucial aspect of which is self-reference in the sense that they try to maintain their own functioning rather than aim at producing an output for somebody *else's* purpose. In the constructivist perspective, modeling living systems—as a procedure to trivialize complex systems in the above sense—must be regarded as turning autopoietic machines into allopoietic ones, i.e., as opening their fundamental closure with respect to the modeler. Maturana notes that "an observer may treat an autopoietic system as if it were an allopoietic one by considering the perturbing agent as input and the changes that the organism undergoes while maintaining its autopoiesis as output. This treatment, however, disregards the organization that defines the organism as a unity by putting it in a context in which a part of it can be defined as an allopoietic subsystem by specifying in it input and output relations" ([22], p. 468).
The conclusion RC suggests is that observed behavior, e.g., defined as protocol of inputs and outputs, cannot capture what a cognitive system knows.

6. CREATING MEANING IN ARTIFACTS

If artifacts should be able to create genuine meaning leading to knowledge, our modeling endeavors must not be guided by copying observable behavior of existing natural systems. Rather, we need a deeper structural insight. The radical constructivist notion of organizational closure is a good point to start with. It is a necessary quality of the nervous system (and hence the cognitive apparatus) of creatures, and is based on the *principle of undifferentiated encoding* of nervous signals. Heinz von Foerster described this ubiquitous neurophysiologic quality as follows. "The response of a nerve cell does *not* encode the physical nature of the agents that caused its response. Encoded is only 'how much' at this point on my body, but not 'what'" ([10], pp. 214–215). In other words, the nervous signals in the young blackbirds that open their beaks at the sight of simple dummies do in no way convey the information of seeing a dummy (or the genuine parent bird it substitutes). Philosophically speaking, the cognitive system is in a brain-in-a-vat situation [33] as it has no independent reference to what has caused the incoming electro-chemical signals. With Maturana and Varela [25], we can compare the situation with that of the navigator in a submarine. He avoids

---

[4] For example, the minus sign corresponds to the action to take something away from something else. This hints at the philosophical framework of "operationalism," as introduced by Percy Bridgman [3].

[5] I follow here the terminology of Heinz von Foerster ([8], p.169) who characterized "anthropomorphizations" as "projecting the image of ourselves into things or functions of things in the outside world."

[6] Cf. also Valentino Braitenberg's [2] "law of uphill analysis and downhill synthesis."

[7] According to Maturana ([22], [24]), autopoietic systems are a subset of self-organizing systems that obey the following criteria: 1. The components of autopoietic systems take part in the recursive production of the network of production of components that produced those components. 2. An entity exists in the space within which the components exist by determining the topology of the network of processes. A system that does not fulfill these criteria is called allopoietic, e.g., machines that serve a different purpose than maintaining their own organization.

reefs and other obstacles without looking even once through the porthole of the vessel. All he needs to do is to maintain a certain (dynamic) relationship between levers and gauges.

Radical constructivism (RC, [15]) is the conceptual framework that builds on this insight. According to the *radical constructivist postulate* [37] the cognitive system (mind) is organizationally closed. It necessarily interacts only with its own states. Or, as Terry Winograd and Fernando Flores [40] put it, the nervous system is "a closed network of interacting neurons such that any change in the state of relative activity of a collection of neurons leads to a change in the state of relative activity of other or the same collection of neurons." Cognition is, therefore, a continuously self-transforming activity. There is no purpose attached to this dynamics, no goals imposed from the outside relative to the cognitive apparatus. It is also in line with the *dreaming machine-argument* of Rudolfo Llinás [19]. Since the nervous system is able to generate sensory experiences of any type, we are facing the fact that "we are basically dreaming machines that construct virtual models…" Llinás's closed-system hypothesis argues that the mind is primarily a self-activating system, "one whose organization is geared toward the generation of intrinsic images." The global picture is that cognition acts independently of the environment. It merely requests confirmation for its ongoing dynamical functioning and works autonomously otherwise: "Although the brain may use the senses to take in the richness of the world, it is not limited by those senses; it is capable of doing what it does without any sensory input whatsoever."

As a result, artifacts with genuine knowledge have to be designed as *closed-loop* systems that regulate their input rather than focus on the production of something different from themselves. As pointed out in the previous section, this definition refers to Maturana's concept of autopoietic systems, which have to be distinguished from allopoietic ones. Autopoiesis refers to mutually chained processes that produce the components necessary to run these processes. Evidently, in the physical space of living systems, autopoiesis is instantiated by material processes, which produce, as it were, material and behavioral by-products visible to an observer. However, these "outputs" do not define the autopoietic system.

The implications of this distinction can be illustrated by comparing a car with a horse. The former needs constant supervision. It is, cognitively speaking, a dumb machine that needs explicit instructions when dealing with its environment, whereas the horse is an autonomous system that continuously modifies its behavior while interacting with the rider. It does not need instructions such as how to deal with obstacles. The car is the product of human technology aiming at creating a machine that can be used for transportation. The horse, however, being the result of evolution, has not been created with a straightforward purpose.

7. KNOWLEDGE IN CLOSED-LOOP SYSTEMS

Evidently, in closed-loop systems, knowledge cannot refer to mapping between an external state of affairs and cognitive structures. The conventional JTB definition of knowledge can no longer be applied in this context since there are neither propositions nor can their truth be specified. In a certain sense, knowledge may be regarded as a belief that receives inductive justification. Following Glasersfeld's characterization of RC, knowledge must not be considered to be passively received but actively built up by the cognizing subject because the "function of cognition is adaptive and serves the organization of the experiential world, not the discovery of ontological reality" ([14], p.182). This leads to an alternative understanding of knowledge that refrains from assuming that differently constructed conceptual frameworks in individuals gradually converge towards an "objectively valid" knowledge system representing the "truth." Since from the perspective of RC no such convergence takes place, the emphasis is to be put on mechanisms of knowledge construction, and on the fact that cognitive systems actively construct their world rather than being passively flooded by information from the outside. Knowledge does not reside somewhere else and is not independent of the cognitive system that generates it. Hence, whatever it is that the cognitive apparatus picks up, it cannot be considered knowledge: "The environment contains no information. The environment is as it is." ([8], p. 189).

This matter of fact is referred to as the *methodological corollary* of RC [37], which says that knowledge is necessarily circular since there is no outside point of reference. This is to be contrasted with the situation in formal systems, in which a set of axioms is introduced *a priori* to serve as a truth criterion. Postulated statements are logically true if they can be derived from the axioms. In these one-dimensional formal systems, e.g., propositional calculus, circularity is a threat. RC, however, locates knowledge embedded in a network of constructive elements that support each other. As such, it is closely related to the coherence model of knowledge that is marked by the quality of "hanging together" rather than being "a helter-skelter collection of conflicting subsystems" [1]. This constructive network in the mind is considered to be without a beginning or ending since its relational structure has developed with a bootstrapping process: Elements are continuously added over time. In the absence of external references, other criteria for checking the validity of a concept must be pulled up, such as coherence, consistency, and richness of referential concepts involved. Coherence means that each element is backed by a number of other elements such that the resulting network of mutual support cannot be flattened into a one-dimensional deductive train of argumentation. In this sense, knowledge is both circular and non-tautological, and emerges in the ongoing dynamics of cognitive processes. Therefore, it is not only processes. Therefore, it is not only *topologically*[8] but also *dynamically distributed*.

Neisser's schemata-guided pickup-paradigm is one way to describe this matter of fact. The configuration of perception-anticipating 'slots' in schemata varies over time, i.e., what is being perceived (and taken in) now may not be perceived at a later instant. This accounts for the variation of meaning we and cognitive systems in general encounter over time. This "constructivist-anticipatory principle" [36] assumes that knowledge is implemented in the form of schemata, which consist of conditions and a sequence of actions. Schemata can be mutually embedded. So can conditions and actions. The purpose of the condition part is to provide context matching which allows the schema that best fits the present context to execute its action sequence. Since conditions can also be part of a sequence, they act as checkpoints for determining whether the anticipated meaning embodied by the schema is still on the right track. After a schema and all its subordinated elements finished, the cycle starts again. In this model, knowledge refers to the capability of the system to bridge between momentary perception and older experiences that are embedded in its schemata.

8. CONCLUSION

Radical constructivism objects to considering knowledge a justified belief that is true in the sense of referentially mapping states of affair in the environment onto cognitive structures. Rather, knowledge must be *system-relative*. This term refers to situated cognitive processes the dynamics of which is merely modulated by their environment on request of the cognitive apparatus rather than instructed by it. Such schemata-guided pickup means *giving up linguistic transparency* [30], which is required to explicitly design systems. Furthermore, it becomes evident that knowledge is a relational dynamical structure rather than a set of propositions. What a person knows today can have a completely different significance tomorrow. This dynamics cannot be captured in a static blueprint we refer to as declarative or procedural knowledge. Rather, knowledge is the process of continuous constructions, the dovetailing of cognitive structures, which occasionally allow for assimilation of and accommodation to picked-up data from the environment. These data (or signals) do not constitute knowledge. Therefore, given this self-referential character of autopoietic systems, knowledge cannot be considered the result of adaptation to an environment.

Future research in computational disciplines will have to focus on the design and implementation of closed-loop systems [39],[9] on how they acquire, represent, and communicate knowledge gained from experiments, as well as on the cognitive and epistemological consequences of using such artifacts.

[8] In the sense of Wittgenstein's "family resemblances" [41].

[9] A recent example of a closed-loop discovery system can be found in [18]. It not only generates and selects hypotheses independent of the human programmer but also carries out the necessary experiments to validate them.

9. REFERENCES

[1] BonJour, L. (1985) *The structure of empirical knowledge*. Cambridge MA: Harvard University Press.

[2] Braitenberg, V. (1984) Vehicles: Experiments in synthetic psychology. Cambridge: MIT Press.

[3] Bridgman, P. (1927) *The logic of modern physics*. New York: MacMillan.

[4] Campbell, Donald. T. (1974) Evolutionary epistemology. In: Schlipp, Paul. A. (ed.) *The philosophy of Karl Popper*. Open Court: La Salle, p. 413–463.

[5] Deneubourg, J.-L., Goss S., Franks N., Sendova-Franks A., Detrain C. & Chretien L. (1991) The dynamic of collective sorting robot-like ants and ant-like robots. In: J.-A. Meyer & S. Wilson (eds.) *Proceedings of the first conference on simulation of adaptive behavior (SAB)*. Cambridge MA: MIT Press, pp. 356–365.

[6] Epstein, J. M. & Axtell, R. (1996) *Growing artificial societies. social science from bottom up*. Cambridge MA: MIT Press.

[7] Feynman, R. (1985) *Surely you're joking, Mr. Feynman!* New York: W. W. Norton.

[8] Foerster, H. von (1970) Thoughts and Notes on Cognition. In: P. Garvin (ed.) *Cognition: A Multiple View*. New York: Spartan Books, pp. 25–48. Reprinted in: Foerster, H. von (2003) *Understanding understanding* New York: Springer-Verlag pp. 169–189.

[9] Foerster, H. von (1972) Perception of the future and the future of perception. *Instructional Science* 1 (1): 31–43. Reprinted in: Foerster, H. von (2003) *Understanding understanding*. New York: Springer, 199–210.

[10] Foerster, H. von (1973) On constructing a reality. In: Preiser, F. E. (ed.) Environmental Design Research Vol. 2. Stroudberg: Dowden, Hutchinson & Ross, pp. 35–46. Reprinted in: Foerster, H. von (2003) *Understanding understanding*. New York: Springer-Verlag, pp. 211–228.

[11] Gettier, E. L. (1963) Is justified true belief knowledge? *Analysis* 23: 121–123.

[12] Glasersfeld, E. von (1974) Piaget and the radical constructivist epistemology. In: Smock, C. D. & Glasersfeld, E. von (eds.) *Epistemology and education*. Athens, GA: Follow Through Publications, pp. 1–24.

[13] Glasersfeld, E. von (1983) Learning as a constructive activity. In: Bergeron, J. C., & Herscovics, N. (eds.) Proceedings of the fifth annual meeting of the North American chapter of the international group for the psychology of mathematics education. Montreal: University of Montreal, pp. 41–69.

[14] Glasersfeld, E. von (1989) Constructivism in education. In: Husen, T. & Postlethwaite, T. N. (eds.) *The international encyclopaedia of education*, 1st edition, Supplementary Volume 1. Oxford: Pergamon Press, pp. 162–163.

[15] Glasersfeld, E. von (1995) Radical constructivism: A way of knowing and learning. London: Falmer Press.

[16] Horgan, J. (1995) From complexity to perplexity. *Scientific American* 272: 74–79.

[17] Keil, F. C. (2003) Folkscience: Coarse interpretations of a complex reality. *Trends in Cognitive Sciences* 7: 368–373.

[18] King, R. D., Whelan, K. E., Jones, F. M., Reiser, P. G. K., Bryant, C. H., Muggleton, S. H., Bell, D. B. & Oliver, S. G. (2004) Functional genomic hypothesis generation and experimentation by a robot scientist. *Nature* 427: 247–252.

[19] Llinás, R. R. (2001) *I of the vortex*. Cambridge MA: MIT Press.

[20] Lorenz, K. Z. (1941) Kants Lehre vom Apriorischen im Lichte gegenwärtiger Biologie. *Blätter für Deutsche Philosophie* 15: 94–125. English translation: Lorenz, K. Z. (1982) Kant's doctrine of the a priori in the light of contemporary biology. In: Plotkin, Henry C. (ed.) *Learning, development and culture*. Chichester: John Wiley, pp.121–143.

[21] Lorenz, K. Z. (1970–1971) *Studies in animal and human behavior*, translated by R. Martin. Cambridge, Massachusetts: Harvard University Press.

[22] Maturana, H. R. (1974) Cognitive strategies. In: Foerster, H. von (ed.) *Cybernetics of cybernetics*. Illinois: University of Illinois.

[23] Maturana, H. R. (2005) The origin and conservation of self-consciousness. *Kybernetes* 34 (1/2): 54–88

[24] Maturana, H. R. & Varela, F. J. (1980) *Autopoiesis and cognition: The realization of the living*. Boston Studies in the Philosophy of Sciences. Vol. 42. Boston, Dordrecht: D. Reidel.

[25] Maturana, H. R. & Varela, F. J. (1987) *The tree of knowledge.* Boston MA: Shambhala.

[26] McAllister, J. W. (2003) Algorithmic randomness in empirical data. *Studies in the History and Philosophy of Science* 34: 633–646.

[27] Nagel, T. (1974) What is it like to be a bat? *Philosophical Review* 83: 435-450.

[28] Neisser, U. (1975) *Cognition and reality.* San Francisco: W. H. Freeman.

[29] O'Regan, J. K. & Noë, A. (2001) What it is like to see: A sensorimotor theory of perceptual experience. *Synthese* 129: 79–103.

[30] Peschl, M. & Riegler, A. (1999) Does representation need reality? In: Riegler, A., Peschl, M. & Stein, A. von (eds.) (1999) *Understanding representation in the cognitive sciences.* New York: Kluwer Academic / Plenum Publishers, pp. 9–17.

[31] Piaget, J. (1937) *La construction du réel chez l'enfant*. Neuchâtel: Délachaux and Niestlé. English translation: Piaget, J. (1954) *The construction of reality in the child*. New York: Ballantine.

[32] Porr, B. & Wörgötter, F. (2005) Inside embodiment. What means embodiment to radical constructivists? *Kybernetes* 34 (1/2): 105–117.

[33] Putnam, H. (1982) *Reason, truth, and history*. Cambridge: Cambridge Univ. Press.

[34] Quine, W.V.O. (1969) Epistemology naturalized. In: Quine, W.V.O., *Ontological relativity and other essays*. New York: Columbia University Press.

[35] Riedl, R. (2000) Strukturen der Komplexität. Eine Morphologie des Erkennens und Erklärens. Heidelberg: Springer-Verlag.

[36] 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.

[37] Riegler, A. (2001) Towards a radical constructivist understanding of science. *Foundations of Science* 6 (1–3): 1–30.

[38] Riegler, A. (2002) When is a cognitive system embodied? *Cognitive Systems Research* (special issue on "Situated and embodied cognition") 3: 339–348.

[39] Riegler, A. (2005) Is a closed-loop discovery system possible? In: Magnani, L. (ed.) *Proceedings of E-CAP 2004*. (to appear)

[40] Winograd, T. & Flores, F. (1986) Understanding computers and cognition: a new foundation for design. Norwood, NJ: Ablex.

[41] Wittgenstein, L. (1999) *Philosophical investigations*, translated by G. E. M. Anscombe. Prentice Hall. Originally published as "Philosophische Untersuchungen" in 1953.