A Symbol Is Not A Symbol

Nils Dahlbäck
Natural Language Processing Laboratory
Department of Computer and Information Science
Linköping University
S-581 83 Linköping, Sweden
nda@ida.liu.se

Abstract

One of the weak points of the present critics of AI is their lack of an alternative cognitive theory. In the paper the outline of such a theory is sketched, based on (one reading) of Piaget's genetic epistemology. This is used to clarify the ambiguity of the terms 'symbol' and 'representation', thereby making it possible to suggest a reconciliation of the positions of the critics and defenders of the representational theory of the mind.

1. Introduction

It is a strange situation. AI is flourishing, scientifically and in an ever growing number of applications. At the same time the field receives critique, not only from without but also from within, the recent book by Winograd & Flores [1986] being one prime example. And the critique is aimed at the very foundations of AI, the knowledge representation hypothesis [Smith, 1982] and the physical symbol system hypothesis [Newell, 1980].

For those cognitive scientists sympathetic to the arguments of the AI-critics, this situation creates an interesting cognitive dissonance, from which at least two interrelated questions emerge: What would a cognitive theory that could accomodate these seemingly contradictory positions look like, and what could such a theory say — or at least indicate — about what computers *can* do?

My aim with the present paper is to show how the two positions of 'mainstream AI' and its critics can be reconciliated by suggesting the outlines of one possible such theory and then suggest what an answer to the second question emerging from this position would look like. I don't really believe that the issue is of utmost importance for the active AI-researcher, who I assume will continue with his work regardless of what philosophically interested scholars will say and do. But it is of some importance for those of us concerned with the impact of AI on society, since this impact comes not only from what existing systems can do, but perhaps to an even larger extent from what the layman and the politicians believe will be possible to do in the future. And this belief is shaped by claims of AI practitioners and critics.

The essence of the argument is simple. I will claim that the two central terms in the debate, 'symbol' and 'representation' are semantically ambiguous, and furthermore that this is overlooked by both sides in the debate, thus making way for an overestimation of the explanatory scope of their theoretical positions. If this claim is accepted, it is then possible to see why critics such as Dreyfus and Winograd & Flores in essence are right in their arguments, but wrong in their predictions.

Nothing of what I will say is really new. What I have tried to do is to compile work from a number of different scholars into a hopefully coherent position, and tried to point out some conclusions which follow from it. Since I am treading in the footsteps of a number of researchers, a large portion of the paper is devoted to a review of these positions. The limited space forces me to keep the reviews short, but I hope I will manage to give a flavor of their respective positions to those not familiar with them (and I am sure I will make readers familiar with the work of Newell, Maturana, Piaget etc shake their heads in despair at my oversimplifications).

2. The knowledge representation hypothesis

Knowledge and *knowledge representation* are important concepts for many areas of cognitive science. For AI, they are central. One could perhaps even say that the question of how knowledge should be represented in intelligent systems, and problems related to this, is the common denominator between the many diverse research areas of AI. For most researchers, knowledge is a symbolic representation, and inference or thinking is done through the manipulation of symbols. This is not something new or unique to AI. In fact, the basic idea was formulated very clearly by Kenneth Craik in the early forties [Craik, 1943], and the information processing paradigm in cognitive psychology shares these assumptions. In AI, Brian Smith [1982] has formulated explicitly the so-called *knowledge representation hypothesis*, which he claims in one form or another lies behind most work in Artificial Intelligence, though never explicitly formulated. He summarizes the hypothesis as follows:

Any mechanically embodied intelligent process will be comprised of structural ingredients that a) we as external observers naturally take to represent a propositional account of the knowledge that the overall process exhibits, and b) independent of such external semantical contribution, play a formal but causal and essential role in engendering the behavior that manifests that knowledge.

The two important points are thus that there exists something that we as observers take as representing 'the knowledge', and that this not only exists but actually is a part of producing the intelligent behavior. The latter part seems uncontroversial, but how is the knowledge represented? The standard answer in AI would seem to be something to the effect that it is symbolically represented. But what then is a symbol? This question is not exactly a new one, it's been with us since the ancient Greeks. But what's interesting in the present context is that Allen Newell [1980] has claimed that a new notion of symbols has evolved within AI. This notion of symbol is intrinsic to the concept of a *physical symbol system*, i.e. a class of systems capable of having and manipulating symbols, and at the same time realizable as a physical system. Newell states explicitly that 'these symbols are in fact the same that we as humans have and use everyday of our lives', and that 'humans are instances of physical symbol systems' [ibid, p 136]. Intelligent programs are physical symbol systems too, this is where we have a connection between human intelligence and the broader class of intelligent systems that encompasses other intelligent machines too. However, he also claims that this notion of symbol 'is distinct from the notion of symbol that has arisen in describing directly human linguistic, artistic and social activities' [ibid, p 141].

I will not here give a detailed account of Newell's rich argument, but concentrate on those aspects that are relevant for the present task, hopefully without doing to much injustice to the original paper. As Newell observes, 'the most important concept for a symbol system is that which gives symbols their symbolic character, i.e. what lets them stand for some entity' [ibid, p 156]. He calls this concept *designation*, claiming that it is more or less synonymous with terms such as reference, denotation, meaning, standing for, etc. He defines designation in the following way:

Designation: An entity X designates an entity Y relative to a process P, if, when P takes X as input, its behavior depends on Y.

This process is closely connected to the process of interpretation, which is defined by Newell as follows:

Interpretation: The act of accepting as input an expression that designates that process, and then performing that process.

I assume that this picture seems familiar to most readers, even for those that have not read Newell's paper, since what Newell in a sense has done, is to have made explicit the meaning and usage of the term *symbol*, as it is used in Artificial Intelligence when for instance talking of LISP programming as *symbolic programming*. For most AI practitioners, this position is obvious and self evident.

However, while this paradigm has dominated, it has not been uncontested. Hubert Dreyfus' *What computers can't do* [Dreyfus, 1972] is a well known, early example. Dreyfus' critique, based on the philosophy of Heidegger has hardly been accepted by the majority of the research community. More recently, Terry Winograd and Fernando Flores have criticized the same tradition from a similar perspective [Winograd, 1981, Winograd & Flores, 1986]. This critique is based not only on Heidegger's philosophy, but also on the biologist Maturana's theories of the organization of living systems and of language and cognition [Maturana, 1975, 1978]. In the next section I shall give a short review of the arguments of Winograd and Flores.

3. Knowledge is not represented

Winograd and Flores challenge the assumption that thinking and other cognitive activities are based on the manipulation of mental representational structures. As previously mentioned, they base their analysis on the work of Maturana, who is a biologist who did his first work on the visual system of the frog [Maturana, *et al.* 1960], where he was led to reformulate his views of the function of the visual system. This work has then led to work on language and cognition, where he has tried to use the results from the work on the visual system to explain these higher cognitive functions.

Starting from the traditional view, in which there exists "an objective (absolute) reality, external to the animal and independent of it (not determined by it) which it could perceive (cognize), and the animal could use the information obtained in this way to compute a behavior adequate to the perceived situation", Maturana came to realize that "the central purpose in the study of color vision could not be the study of a mapping of a colorful world on the nervous system, but rather that it had to be the understanding of the participation of the retina (or nervous system) in the generation of the color space of the observer". [Maturana, 1970]

The view that we don't have direct access to the outside world, but that our experience is partly determined by the activities in the nervous system, is not unique these days. But Maturana goes one step further when he claims that the cognitive system does not make *any* use of any representational structures. The light striking the retina triggers chemical changes in the neurons, which causes the structure of the nervous system to change (cf the neurological argument in sec 6). The arguments by Winograd & Flores are obviously more intricate than this. But I hope the foregoing sections make it possible to get a glimpse of a central part of their argument, namely that it is not only possible to explain intelligent behavior without using symbolic representations, but that any attempt to do so is based on a faulty epistemology, separating organism from the environment. (The tradition of doing so is called by them *the rationalistic tradition*. There seems to be a close resemblance between this and Lakoff's [1987] concept of the objectivist paradigm).

4. What happened to cognition?

Some cognitive scientists might dismiss the arguments by Winograd & Flores, on the grounds that they only consider biological (Maturana) and philosophical (Heidegger) aspects, and that therefore their arguments apply to other domains of explanation than cognition. Maturana talks about cognition, but in a sense quite different from the way the term is usually used. One could argue that before the critics of the representational theory of the mind have formulated an alternative *cognitive* theory, their position has little to offer in way of an alternative for the cognitive scientist, be he psychologist or AI researcher. (There exists of course an alternative to this, and that is to claim that we don't need any cognitive theory to explain the mind. This seems to be the stance taken by Searle [1984]).

However, in my opinion there already exists such a cognitive theory, or at least the skeleton of such a theory, and that is the genetic epistemology of Jean Piaget. I furthermore believe that this theory can be used, not only to fill the gap in the aforementioned critique, but also to clarify some of the issues raised in the debate between critics and defenders of the representational theory of the mind.

5. A biological epistemology

Here I will not try to give any general overview of Piaget's work, but instead confine myself to an introduction to his theory of knowledge, and the view of symbols and representation which follow from that theory. It is important to realize that what follows is different on central points from most English and American expositions of Piaget. I am of course not claiming that my reading is the only "correct" reading of Piaget's work. Anyone who has read him in the original will realize how outrageous such a claim would be. Piaget is not exactly an easy writer! And I am not interested in some exegetic analysis of the numerous books and articles published. My aim is simply to summarize one plausible reading, which I would claim is interesting in and of itself, and which also has interesting consequences for a discussion on the nature of knowledge and representation.

Jean Piaget is best known as a child psychologist. But what makes him important is not perhaps primarily what he has taught us about the mental development of the child, but his use of experimental psychological methods, used within a biological framework, for studying epistemological questions. In his own words: "My most central concern has always been to determine the contributions of the person's activities and the limiting aspects of the object in the process of acquiring knowledge" [Flavell, 1963, p vii].

For Piaget, a biological organism implies a structure which is responsive to its environment. The stimulus is never something which is "out there", but rather that aspect of the environment to which the organism is responding. The central factor in this is the underlying structure of the organism. The stimulus is assimilated to the underlying structure, and if it could not be assimilated, it would simply not exist for that organism.

This structuring exists on all levels of behavior. Knowledge is used as a general term, much wider than conscious knowledge, and is in fact synonymous with this structuring aspect of action. Knowing is therefore an activity of the subject, and knowledge is something constructed. That this view of knowledge as action-knowledge is relevant for lower organisms is not perhaps unique. It is has also been a central part of von Bertalanffy's [1973] General System Theory, which in turn was inspired by von Uexkull's *Umwelt-Lehre* [von Uexkull, 1920] which was summarized by von Bertalanffy as follows: "From the great cake of reality, every living organism cuts a slice, which it can perceive and to which it can react, owing to its psycho-physical organization, i.e. the structure of receptor and effector organs." [von Bertalanffy, 1973, p 240] This position seems also to be similar to that of Maturana. But what perhaps makes Piaget unique is his claim that *all* intelligence and knowledge can be and should be seen in this perspective: in his insistence that knowledge is not caused by nor upheld through symbolic representations, while at the same time giving place for the symbolic or representational activities of the mind within an integrated theory.

Piaget distinguishes between two forms or types of knowledge, *operational* or *operative*, which is based on a logico-mathematical type of activity, of ordering, counting, classifying, i.e. activities for which the objects as such are no more than a support, and *figurative*, which is aimed at extracting information from the objects themselves, such as their colors, form, weight etc [Inhelder, 1962]. (There is an obvious similarity between Piaget's distinction operative-figurative knowledge and Tulving's [1972] distinction between semantic and episodic knowledge). Piaget has primarily studied the operational knowledge - in a way you can say that he has studied the development of Kant's categories.

The operative knowledge, the logico-mathematical knowledge, is for Piaget the primary knowledge: And this in two senses; it is considered the central aspect of thinking, with the figurative in some sense secondary or dependent on it, and it is the most studied aspect of the intellectual development. Consistent with his view of knowledge as action-knowledge, an operation is an action in the same literal sense in which making a detour to find a desirable object is an action. [Furth, 1967] But it is not the action as such, but the generalizable aspect of the act which is central, and which also is called a *scheme* ("A scheme is, in fact, that which is generalizable in a given action" [Piaget, 1963]). These structures or categories are not inborn and ready-made, but the result of a developmental process, and the knowledge is not an internal copy or representation of the environment, but the action-structures which guides the organism-environment interaction. And these action-structures are not symbols. As Furth puts it: "For Piaget representation and symbol are consequences of knowing, not explanatory antecedents or intermediaries" [Furth, 1967, p 20]. More on this later.

Piaget uses the term *figurative knowledge* to refer to the static, figural and particular aspect of knowledge. He contends that this aspect is dependent on the operative knowledge, in the sense that perception, imagery etc is dependent on the current level of operative development. (To give one example of this dependency: before the development of eye-hand coordination, which is crucial for the development of the scheme of the permanent object, the child is unable to choose the bigger of two objects, when the bigger is further away and therefore has a smaller apparent size on the retina.) I will not go deeper into this part of Piaget's theory here for two reasons. It is not needed for the rest of the argument, and it is a part of the theory that clearly is in need of further development.

Instead I want to point out that there is an obvious similarity between Piaget's and Maturana's positions. What seems to be a common denominator is the realization that the organism never reacts to its environment as a 'tabula rasa', but only through its perceptual and cognitive structures, and that it is these as much as the 'objective world' which determines what is and is not the 'umwelt' of the organism.

A central question in this context is obviously what makes it possible for an organism to react adaptively to the world. Piaget uses a 'pragmatic' concept to explain this. He claims that there is a tendency for equilibration between organism and environment, and within the operative structures. In effect, this is only two sides of the same coin. As with other aspects of Piaget's theory, the concept of equilibrium is not easy to understand. In fact, most of his readers would probably claim exactly the opposite! (For a discussion of this and other aspects of Piaget's theory, linking it also to AI, see Boden, 1979) But for our present purposes it is sufficient to note that Piaget has started to develop a theory which makes it possible to explain the development of adaptive reactions to the environment without having to postulate some sort of 'copy--knowledge' of the world as it is. Another important point is the emphasize on viewing the organism in relation to its environment, which points to the close connection between knowledge and adaption (In a sense, Newell's [1982] *The Knowledge Level* expresses a similar point from a different perspective).

5.1 Representation

I claimed earlier that Piaget explains knowledge without the use of symbolic representations, while at the same time giving place for symbolic activities within the theory. To understand this, we need to clarify the meaning, or rather meanings, of the term 'representation'. The following discussion leans heavily on the work by Furth [1967].

The term 'representation' can be used in both an active and a passive sense. Furth claims the active sense to be the primary one "to make something present by means of ..". Here the person is the subject of the activity, and the symbol is a mediating instrument. In the passive sense the person has "disappeared", and the symbol becomes the subject of the sentence, as in "this map represents the city of Timbuktu". In this case there is some figural resemblance between the symbol and that which it represents, but there are also cases where the relation is more abstract or conventional, words in natural language being an obvious example.

In modern cognitive science, with one of its roots in 17th century empiricist philosophy, the latter meaning of representation is the most common one. In this tradition, the internal representation of external reality is the chief explanatory factor for intelligent behavior, and knowledge is more or less synonymous with symbolic representations.

In contradistinction to this, Piaget retains the active use of the term representation, and explains knowledge without representations, as was pointed out in the previous sections. Or, perhaps it would be more correct to say that he is aware of the two uses of the term, and in discussing symbolic behavior proper (language, images, dreams etc.) keeps this distinction in mind

In fact the word 'representation' is used in two different senses. In the wide sense, representation is identical with thought, that is, with all intelligence which is not simply based on perceptions or movements (sensory-motor intelligence), but on a system of concepts or mental schemes. In the narrow sense, repre-

sentation can be limited to the mental image or to the memory-image, that is to the symbolic evocation of absent realities. Moreover, it is clear that these two kinds of representation, wide and narrow, are related to each other insofar as the concept is an abstract scheme and the image a concrete symbol; even though one no longer reduces thought to a system of images, it is conceivable that all thought is accompanied by images. For if thinking consists in relating significations, the image would be a 'signifier' and the concept a 'significate'. *[Piaget, 1946, p. 68]*

I will later try to show that this distinction between two meanings of 'representation' is paralleled by a necessary distinction between two meanings of the word 'symbol'.

5.2 Symbols

Thus, for Piaget the use of symbols is always seen in relation to the active aspect of representation mentioned before. They are seen as re-presentations, i.e. the evocation of absent realities. Central to all symbols is the differentiation between the sign and its signification (which is something more than its referent). It is important to notice that Piaget's denial of the symbolic nature of knowledge is closely tied to his notion of symbols, where the user of the symbol has the capacity to differentiate between the symbol and the symbolized, and that his comments on this issue should bee seen in the light of his arguments against an epistemological position that following Hume makes no fundamental distinction between symbols such as images and objects of knowing.

A critical point is of course what an operative structure (or a scheme) is. Wason & Johnson-Laird [1972] have shown that Piaget's view of the so-called formal abstraction is not correct. There is no place in the present context to even hint at a way to resolve the problem. For those readers familiar with the work of Wason & Johnson-Laird I just want to point out that I think the critique can be handled without damaging those aspects of the theory of concern for us now (sorry for the hand-waving!).

6. 'A symbol is not a symbol'

So, after this long, but still very condensed and difficult to digest exposition of some central aspects of Piaget's theory, I will try to use it for a hopefully clarifying discussion of the twin concepts 'symbol' and 'representation'. And this can be done without having to accept Piaget's theory in all its details.

As we saw in the beginning, critics of AI and cognitive science in general, such as Dreyfus and Winograd & Flores, have as one of their favorite arguments that the mind is not symbolic in nature. And on the other side we have the Knowledge Representation Hypothesis and the Physical Symbol System Hypothesis, explicitly formulated by Smith and Newell, but more or less implicitly accepted by the majority of workers in these fields, and whose major point is that the mind *is* symbolic in nature, and that the development of artificial symbolic manipulation devices, i.e. computers, makes the development of artificial intelligences possible. And my simple point is, that in some sense both are correct! Correct in their premises that is, but due to a lack of understanding of the limitations of their theory's domains of explanation, wrong in their conclusions.

If by a symbol we mean something which makes it possible to make an absent reality present in the same way that a child uses a shoe box as a symbol of a bed, i.e. with a knowledge on part of the symbol user that the symbol *is not* that which it represents, then it is obvious that an AI system or Newell's Physical Symbol System is not a symbol-using system in this sense, and the 'knowledge representation' is not a representation in the narrow sense of the word. The 'symbols' are simply something for the program to react to. But they stand for nothing, they mean nothing to the program. An analogy from neurology can perhaps clarify this position for the doubting reader.

Penfield & Roberts [1959] have shown that electrical stimulation of neurons in the brain of patients undergoing brain surgery with local anaesthesia caused visual experiences in these patients. My claim is that the action of the neurons reacting to this simulation are essentially similar to a running computer program. The input is something to react to automatically, and it makes no sense to ask what it means to the system. It has no meaning, it is not a symbol in the strict sense of the word, and, as Penfield and Roberts have shown, it makes no difference for a local point in the system where the 'stimulation' comes from. It's there, and the reaction is determined by this and only this (if the structure of the neuron or LISP function or whatever is taken as given). This is related to Searle's often quoted claim that the ('symbolic') programs 'have only syntax but no semantics' [Searle, 1980].

It is in *this sense* that the mind is not symbolic in nature, and knowledge is not represented, and it is this that I take as the core aspect of and common denominator of the theories of Piaget and Maturana. And if we take knowledge to be the generalizable aspect of the organism-environment interactions, there are no computers that have 'knowledge'. And it is here that I think that the positions of the AI/Cognitive Science critics are essentially correct.

But it seems to me that they have missed an important point. Notwithstanding their criticism of current theories in the field, humans *can* use symbols! And use them, also in thinking. We use an inner language for solving problems, we use imagery for problem solving and daydreaming; symbols function as a support for a number of cognitive activities, and sometimes the way we have represented (in the strict or active sense) the problem can actually prevent us from finding the solution.

7. A synthesis

Let us then try to formulate a synthesis which can accommodate both uses of the term 'symbol'. It would look something like this: There exists two different levels or domains. The first, and in some sense primary one, is the organism/environment interaction as described by Maturana, or the sensory-motor and operational knowledge of Piaget, and for those of you familiar with his theory, it also resembles the Cartesian automata of Johnson-Laird [1983]. Behavior can on this level be described and explained without postulating any mental symbolic representations in the strict sense.

The second level is the symbolic level proper. This exists as a superstructure, as it were, in some organisms — at least in humans. They can use symbols, especially language and images to reflect, plan etc; i.e. to reason about future actions in 'mental models' in their working memory. This has all the disadvantages pointed out by Winograd & Flores when the gap between the symbolization and the actual is too wide. But it is still of immense value in well understood but complex domains. Especially if the agent is aware of the fact that the model is only a model, and therefore is alert to the possibility that it can be incorrect. After all, it is only a map, and maps should not be confused with the reality, as we all know.

8. And the consequences for AI?

The situation for AI is somewhat paradoxical. As pointed out previously, the computer system cannot be seen as a symbol-using system in the strict sense. Some readers might then draw the conclusion that this fact holds a promise that sooner or later we will be able to create intelligent machines, with an intelligence similar to the human intelligence, since I also claimed that this is what characterizes it too. But I would claim that the only possible conclusion to draw is in fact exactly the opposite! And here is the reason:

Since the computer is not a biological organism, interacting with the environment, its 'knowledge' has to be spoon fed, as it were, into the system. And the person doing the spoon-feeding is the programmer. And the only knowledge that *he* can use for this process, is knowledge that for him is symbolized. Symbolization is a prerequisite for communication of complex ideas or concepts. In a sense, the computer and its program can be regarded as an externalization of a mental symbolic model.

This explains why the most successful intelligent programs (regardless of their use of AI techniques or not) has been developed in areas where there exists a usable symbolization of the necessary knowledge. Chess programs are a good example of this. The necessary knowledge had been symbolized and developed through centuries to make it possible for chess players to communicate their experiences. And the 'only' thing the the program developers had to do was to adapt this to the strengths and weaknesses of the computer. And of course it works, as do a number of 'number crunching' programs, for which the situation is essentially the same. And they work even better than humans, since the computer is not hampered by the processing limitations of man's conscious mental processes such as slow speed and memory limitations.

The critique of Dreyfus & Dreyfus that this is not the way a human functions is then almost correct, but not in one essential aspect; When a person needs to consciously manipulate his knowledge, or communicate it, *then* he uses symbolic representations. And these representations can be manipulated in a computer. Therefore, we can say that there are two plausible answers to the question if an AI program is symbolic. Seen from within the program, it is not symbolic in the strict sense. But for us, the users of the system, it manipulates symbols. It is an extension of our working memory.

So the key issue when discussing what expert systems will be able to do in the future is not if the mind uses symbolic representations or not. In most cases it doesn't. But that is not the point. Consequently, the success of expert systems in this perspective will not be dependent on the knowledge engineers success in uncovering the rules and representations that the expert uses. Even if Dreyfus & Dreyfus [1986] have almost exclusively studied a domain (chess) whose rules can be explicitly formulated and where clearcut criteria for success can be established — something which makes the generalizability of their findings an open question — we can accept their claim that experts don't use rules.

The key issue is instead whether we can find ways of symbolizing and representing knowledge that before the advent of AI systems we had no need to handle in such ways.

A corollary of the position put forward here is that even if the critiques of the representational theory of the mind are essentially correct, in that the mind is not a symbolic system, there is no way to mark in advance the demarcation line between what can and cannot be handled in an AI-system. This will all depend on the possibility of developing formal symbolic representations for knowledge where such does not presently exist, simply because previously there was no need for them. Non-monotonic logics seem to be a prime example of this.

9. Much ado about nothing?

So what was all this about? The consequences for AI don't seem all that startling. Just continue to develop formalisms for representing and manipulating areas of knowledge not previously manipulated. It is an open question how successful we will be, and this question is in a sense empirical rather than theoretical. Claiming otherwise would be as sensible as claiming that Zenon was right when he said that Achilles would never catch up with the turtle, just because at that time there did not exist an efficient formal way of representing the problem.

But, as I said in the beginning, the importance of arguments such as these lies not so much in their consequences for the AI reserarcher in his work, but rather — if they are accepted — in modifying some claims about what computers can do and not do in the future. So philosophical arguments are more important for their social implications than for their scientific consequences.

On the other hand, there *are* some important consequences for another branch of cognitive science, namely psychology. But that, as Hans Christian Andersen said, is another story.

Acknowledgements

Many friends at the Department of Computer Science, Linköping University, and Department of Psychology, Stockholm University have participated in discussions related to the issues in this paper. None of them had a chance to influence the content of the paper, so not to embarrass anyone I just thank you all for enjoyable and sometimes clarifying discussions. Special thanks to Ivan Rankin for his effort to transform my Swenglish into English.

References

[von Bertalanffy,1973] Ludwig von Bertalanffy. *General Systems Theory*, Harmondsworth: Penguin.

[Boden,1979] Margaret A. Boden. *Piaget*, Glasgow: Fontana

[Craik,1943] Kenneth J.W. Craik. *The Nature of Explanation.* Cambridge: Cambridge University Press.

[Dreyfus 1972] Hubert L. Dreyfus. *What Computers Can't Do.* New York: Harper & Row.

[Dreyfus and Dreyfus 1986] Hubert L. Dreufus and Stuart E. Dreufus. *Mind Over Machine* Oxford: Blackwell.

[Flavell 1963] John H. Flavell. *The Developmental Psychology of Jean Piaget,* New York: Van Nostrand Reinhold Company .

[Furth 1967] Hans G. Furth. *Piaget and Knowledge: Theoretical Foundations*, Engelwood Cliffs, N.J.:Prentice-Hall

[Inhelder 1962] Bärbel Inhelder. Some Aspects of Piaget's Genetic Approach to Cognition. In: W. Kessen & C. Kuhlman (eds.) *Thought in the Young Child. Monographs of the Society for Research in Child Development,* 27, 19-32.

[Johnson-Laird 1983] Philp N. Johnson-Laird. *Mental Models* Cambridge: Cambridge University Press.

[Lakoff 1987] George Lakoff. *Women, Fire, and Dangerous Things: What Categories Reveal About the Mind*, Chicago: Chicago University Press.

[Maturana 1970] Humberto R. Maturana. Neurophysiology of cognition. In P.Garvin(ed.) *Cognition: A Multiple view*, New York: Spartan Books.

[Maturana 1975] Humberto R. Maturana. The organization of the living: A theory of the living organization. *International Journal of Man-Machine Studies,* 7, 313-332.

[Maturana 1978] Humberto R. Maturana. Biology of language: The epistemology of reality. In G.A. Miller & E. Lenneberg (eds.) *Psychology and Biology of Language and Thought: Essays in Honor of Eric Lenneberg,* New York: Academic Press.

[Maturana *et al.*, 1960] Humerto H. Maturana, J.V. Lettvin, W.S. McCulloch, and W.H. Pitts. Anatomy and physiology of vision in the frog. *Journal of General Physiology*, 43, 129-175.

[Newell, 1980] Allen Newell. Physical Symbol Systems, *Cognitive Science*, 4, 135-183.

[Newell, 1982] Allen Newell. The Knowledge Level, *Artificial Intelligence*, 18:87-127.

[Penfield and Roberts, 1959] W. Penfield and L. Roberts. *Speech and Brain-Mechanisms* Princeton: Princeton University.

[Piaget 1946] Jean Piaget. *La Formation du symbole chez l'enfant*, Neucahtel: Delachaux et Niestle. (*Play, Dreams and Imitation in childhood* Norton, 1951)

[Piaget 1963] Jean Piaget. Le langage et les operations intellectuelles. In: *Problem de psycholinguistique*, pages 51-61, Paris: Presses Universitaires de France. (English translation in Furth[1967]).

[Searle 1980] John R. Searle. Minds, Brains, and Programs. *The Behavioral and Brain Sciences*, 3, 417-424.

[Smith 1982] Brian C. Smith. *Reflection and Semantics in a Procedural Language* Tech.Report MIT/LCS/TR-272. MIT, Cambridge, MA.

[Tulving 1972] Endel Tulving. Episodic and Semantic Memory. In E. Tulving & W. Donaldson, editors. *Organization of Semantic memory.* New York: Academic Press.

[Wason and Johnson-Laird 1972] Peter C. Wason and Philip N. Johnson-Laird. *The Psychology of Reasoning: Structure and Content* London:Batsford.

[von Uexkull 1920] J. von Uexkull. *Umwelt und Innenwelt der Tiere* 2nd ed. Berlin: Springer.

[Winograd 1981] Terry Winograd. What does it mean to understand language? *Cognitive Science*, 4, 209-241.

[Winograd and Flores 1986] Terry Winograd and Fernando Flores. *Understanding Computers and Cognition*, Norwood, N.J. Ablex.