How situated is cognition?

Jacobijn Sandberg and Bob Wielinga
Social Science Informatics
University of Amsterdam
Roeterstraat 15
1018 WB Amsterdam
The Netherlands
Telephone: (31) 20 5256789
email: sand berg® swi.psy.uva.nl

Abstract

This paper discusses the recent views on knowledge, representations, and memory as presented by different researchers under the flag of 'situated cognition*. The situated view implies a radical shift of paradigm. We argue that there are no strong reasons to leave the traditional paradigm of cognitive science and AI. Four main issues are addressed; the role of computational models in theories of cognition, theories on knowledge and memory, the frame of reference problem and implications for learning and instruction. The main conclusion of the paper is that 'situationism' is throwing out the baby with the bath water. Consolidated achievements of Cognitive Science and AI still stand, even if the architectures that are assumed to underly traditional models of cognition can be challenged.

1 Introduction

The last few years the term "situated*[1] appears more and more in articles on the subject of learning and teaching [Brown *et a*/., 1989]- [Suchman, 1987];[Harel and Papert, 1990];[Pea, in press]., The use of this term reflects a major shift in thinking about "knowledge", "information", "representation", and even "memory". This rethinking of basic terms related to learning and teaching is not limited to the educational field but has recently penetrated the field of AI [Clancey, in press];[Claneey, 1990b]. What does this shift entail and what is the importance of the concept "situated"? In short, the shift states that knowledge can no longer be viewed as a self-contained substance. Knowledge is inherently indexical of the environment, it can not be separated from its context, knowledge is essentially *situated.* Knowledge is not objective but subjective, embedded in a particular frame of reference. Knowledge is always relative, and open to reinterpretation. For the educators this new view on knowledge has major educational implications. They see learning as a process of enculturation. By joining a community of practice, whatever the practice may be, one gets to know the culture of the community, the jargon used, the beliefs held, the problems raised, the solving methods used. The educators claim that involvement in a community of practice is the only way to learn [Brown *et ai,* 1989]. Formal schooling does present students with a culture, but this school culture has nothing to do with the cultures that surround the subjects to be taught. Thus, students are engaged in tasks and endeavours that have no authenticity, no real-life value. Students do not learn the subjects themselves, but they learn *about* subjects [Brown *et* al., 1989] No wonder, the educators say, that formal education fails. It is no use to try to transfer isolated, decontextualised bits of knowledge. Knowledge can only be gained in authentic activity, it is not a substance, but gets constructed in action. If you want your students to learn mathematics you have to give them the opportunity to act as mathematicians.

This is the situated view on education. How does this view relate to Artificial Intelligence? Work in the field of autonomous agents has clearly adopted a situated view by stressing the interaction of the agent with its environment [Rosenschein, 1985];[Agre, 1988]; [Brooks, 1991). More recently, the situated view has been taken much further by Clancey, by adopting it as a framework for an architecture for cognition and AI at large. He claims that his proposed framework dramatically changes our views on knowledge, representations and memory. He argues, in line with the educators, that knowledge is not a substance, and following from there that representations are not structures in the minds of people, but external, perceivable structures, open to debate, negotiation and reinterpretation. People create representations in their actions. For instance, in speech we do not translate a stored structure into words but construct a new representation each time we speak. Memory can no longer be seen as a storage place for representations - there are no representations left to store - but in Clancey's view should be considered "a capacity for creating and recombining phrases of behavior" [Clancey, in press, page 60].

Thus the representations we equip, for example, our expert systems with do not reflect structures that cause expert behaviour, but are our interpretations of expert behaviour over time. And the expert system does not use or view these representations in the same way people do. To the system these representations are structures to be manipulated syntactically, not semantically. We as observers interpret them semantically, the system however cannot.

In this article we investigate whether the situated view really implies a radical change in our thinking on knowledge, representations and memory and to what extent the suggested implications for learning and teaching are supported by theory and empirical findings. Since the literature on situated cognition is diverse and not all issues raised can be addressed in a short paper, we will focus on four main issues: the role of computational models in the study of cognition, theories of knowledge and memory, the frame of reference problem, and the implications for learning and instruction,

2 The Identity Hypothesis

An important objection that situationists from different camps raise against traditional views on cognition is that intelligent behaviour should not be seen as driven by a knowledge base stored in memory. Knowledge is not to be viewed as stored structures similar to the knowledge representations that we embody in AI programs [Clancey, in press]. In a similar spirit, [Brown *et al.,* 1989] reject the separation between knowledge and doing, which -they claim- treats knowledge as an "integral, self-sufficient substance". They blame traditional education for viewing instruction as the transfer of declarative structures into the head of the pupil.

The situationists' attack on traditional cognitive and instructional sciences, however, starts from an incorrect assumption. This assumption is what we call the *identity hypothesis,* which states that the information and knowledge that we manipulate and represent in our computer programs and teach in our schools are identical or at least very similar to whatever people have in their brains. This assumption does not do justice to much work in cognitive psychology and cognitive science. Cognitive science is building *models* in terms of computer programs or otherwise, which are *abstractions* from the real thing. Such models generate predictions in terms of behaviour, which can be tested, and as such they can be seen as instances of theories of cognition, but only at a certain level of abstraction. An expert system that solves problems through behaviour similar to that of a human expert, can be viewed as a theory of the problem solving behaviour of that expert. However, good fit between model and data does not necessarily mean that all underlying machinery (representations, processes, hardwetware) is identical. The identity hypothesis is a composite one in the sense that it covers different levels of abstraction. The identity hypothesis not only implies an identity relation at the level of observable behaviour, but also at the level of the structures that are being manipulated by certain processes, and at the level of the mechanisms that make it possible to operate on these structures. We think that the identity-hypothesis was never widely held in this strict sense. What might then be the communal view on the relation between human thinking and machine operating as provided by cognitive theory? We view this relation in terms of what we call the *functional-equivalence hypothesis.* To illustrate this point we cite Dennett [L981]: "When an AI model of some cognitive phenomenon is proposed, the model is describable at many different levels, from the most global phenomenological level at which the behaviour is described (with some presumptuousness) in ordinary mentalistic terms down through various levels of implementation, all the way down to the level of program code - and even further down, to the level of fundamental hardware operations, if anyone cares. No one supposes that the model maps onto the process of psychology and biology all the way down". This last sentence is crucial, for it rejects the identity hypothesis. Dennett continues by saying: "The claim is only that for some high level or levels below the phenomenological level (which merely sets the problem) there is a mapping of model features onto what is being modelled: the cognitive processes in living creatures, human or otherwise. It is understood that all implementation details below the level of intended modelling will consist of cognitive wheels - bits of unbiological computer activity mimicking the gross effects of cognitive subcomponents by using methods utterly unlike the methods still to be discovered in the brain'*. We quote Dennett at length to clarify the issue at hand. Most cognitive theories are high-level theories. They are not committed to statements about lower levels. However, to be implemented, the lower levels have to be there as well. That is where the cognitive wheels introduced by Dennett come in. They are purely technical devices, solely there to make the program work. They are not intended or assumed to bear any relation to what might be found in the human brain. However, at the higher level at which the statements of a cognitive theory lie, functional equivalence is hypothesized. The knowledge structures we describe, the representations we form, are assumed to be functionally equivalent to whatever devices humans use to solve problems, in the sense that they give rise to similar behaviour. The functional-equivalence hypothesis reflects the view generally held by the AI community better than the identity hypothesis as assumed and consequently attacked by Clancey in press. A further reason to refrain from any form of identity hypothesis is clearly stated by Anderson [1978, 1979], He has pointed out that there are problems with the use of behavioural criteria in deciding upon the psychological validity of a particular computational representation. Different computational representations (imagery representations versus propositional representations, for example) can in principle predict the same behaviour. There are no objective criteria - as yet - to favour one type of representation over another.

In the field of education, the identity hypothesis crops up in critiques of traditional instructional methods. The essence of this criticism is that instruction in the school context is aimed too much at the transfer of knowledge in declarative, decontextualised form, out of context, and disconnected from its use. However, instructional psychology has been well aware of the problems with the teaching of declarative knowledge [Simon, 1980]. Many studies, e.g. in the fields of mathematics and physics problem solving, have shown that just learning the declarative subject matter is insufficient for operational use of the acquired knowledge. The main problem however has always been how to design teaching methods that both teach the declarative subject matter and its *use* [Mettea and Pilot, 1980] in a representative set of contexts. The problem then is how to choose a set that facilitates generalisation of use of the declarative subject matter. In any case, the assumption that traditional instruction is solely based on the identity hypothesis, is an oversimplification.

In summary, the situationists' reaction against traditional views on cognition, is overstated, because it attacks a position not held by many. Many cognitive scientists and educationers have been well aware of the limitations of their theories and methods.

3 Knowledge and Memory

The situated view on cognition leads Clancey to the view that knowledge is not something that we store as structures in memory. Clancey calls this the - *nalisaiion move.* Through this externaiisation Clancey places knowledge (as understood at the knowledge level [Newell, 1982] as well as symbolic representations *outside* the thinking agent, Knowledge and its symbolic representations are the results of a sense-making process in which an observer describes patterns of behaviour of an intelligent agent.

Clancey then proceeds with the description of a functional architecture which will support this view and in which perception plays a key role. The central role attributed to perception is crucial to the externaiisation move. According to Clancey, we can only know what we can perceive. Constructing a representation means seeing something in a new light. Each act of speaking is a complete act of perceiving in itself By speaking we create new meanings which are perceivable by ourselves and others and thus open to reinterpretation.
Memory, i.e. what is retained from previous activity, is a capability for replaying previously enacted sequences of behaviour ("phrases"). New phrases can be constructed through substitution of actions and recombination of phrases.

How, one may ask, does the functional architecture that Clancey proposes relate to classical theory and the phenomena it tries to explain? Classical theory offers explanations for, amongst others, phenomena of recall, expertise and recognition. Does Clancey's functional architecture explain these phenomena any better?

3.1 Recall

First, there is the obvious observation that people can recall things. We can recall names, numbers, sentences from a poem, laws of physics and so on. Although empirical evidence clearly shows that recall is much better in contexts similar to the context where the information was acquired, recall of decontextualised information is possible and even desirable in problem solving. Remembering Newton's law in the form F—ma does not remind us of all the physical situations in which we have applied this law, nor of the teachers and professors that have explained the law to us. So, how does Clancey's functional architecture explain recall of decontextualised information? We have to assume that some process is replayed that generates the symbolic representation of Newton's law. If such a memory capability exists, why not call it a piece of knowledge? Why can we not create such a memory phrase through explicit communication (i.e. teaching decontextualised subject matter)? Even if in the neural machinery that implements the functional architecture there is no recognisable *place* where a behavioural phrase is stored, at least functionally it can be viewed as a unit that can be activated as such.

3.2 Expertise

Recent experiments on the nature of expertise, show that it is possible to teach new strategies to experts which normally use other, sub-optimal procedures [Staszewski, 1988]. Such instruction not only changes the expert's behaviour, but also enhances performance. In terms of Clancey's functional architecture, this means that process memory *is* extended with new phrases through the perception of external representations (symbols explaining a new procedure), not through actual acting. It is not clear to us whether this corresponds to the reflection process postulated by Clancey, but in any case the functionality of the classical memory model can explain this phenomenon just as well as the process memory architecture, if not better.

A second point concerning expertise is that there is ample evidence that experts in various fields have the capability of recognising large amounts of patterns that can occur in problem situations [de Groot, 1966]; [Ericsson and Poison, 1988]; [Chase and Ericsson, 1981]. Such patterns can be identified empirically and can sometimes be verbalised. Admittedly such patterns are primarily indexical and functional. They impose structure on the observed reality and through that structure make a focussed problem solving process possible. However, there appears to be no good reason why such patterns can not be viewed as a psychological reality. The fact that an observer or an experimenter interprets eye-movement data or verbal utterances of an expert in terms of knowledge structures, and that consequently statements about that knowledge are subjective, relative to the observer, does not change the fact that such patterns can be induced by explicit instruction, or can be elicited in knowledge acquisition. Again there does not appear to be a compelling reason why a process memory should give a better account of these phenomena than a classical theory of memory.

3.3 Constructive memory

A third point concerns the observation that memory is constructive and dynamic. Knowledge evolves continually as it is being used [Brown *ei* al., 1989, page 33]. Indeed there is ample evidence that memory is (re-) constructive. Not only Bartlett's [1977] seminal work on remembering provides strong evidence of this, but many, more experiments support this. For example, the experiments by Bransford nd Franks [1971] and other, similar experiments, clearly indicate that what is remembered is not often the exact stimulus, but the gist, the meaning of a linguistic utterance. In reconstructing the original stimulus, inferences are made and other knowledge is used.

Again, the process memory does not seem to give us much more than conventional theories of memory storage, memory retrieval and depth of processing. The depth of processing of a sentence will determine the manner in which it is stored in semantic memory, and will determine to what extent it can be reproduced or recognised. The deeper the processing, the more decontextualised the information that is stored will be.

Thus, even though many of the observations about knowledge - implied by the functional architecture proposed by Clancey - e.g. its indexical nature, better recall in context etc., are relevant and consistent with empirical evidence, there appears to be no compelling need for a radical change in the paradigm of AI with respect to the notions of knowledge and memory. Of course, it is more likely that an architecture based on neural processes resembles what goes on in our heads than a von Neuman architecture. But that is not the point. As Marr [1977] has pointed out, we are studying information processing *problems* and not in the first place mechanisms.

As to the issue whether knowledge can be viewed as a substance, all evidence cited above strongly suggests the conclusion that knowledge can still be viewed at the functional level as "mental substance" more or less of which can be available to a thinker. Moreover, this substance has internal structure which can be made explicit in empirical studies. Whether this "substance" is stored as symbolic structures, as states of neural nets or as process memories, as Clancey suggests, is immaterial for the present discussion, although it is of course an important question from the architectural point of view. What is our concern here, however, is the evidence that knowledge can at least functionally be viewed as a substance that can be communicated in such a way that it can be used.

4 The frame of reference problem

The externalization move [Clancey, in press) as described above is related to what Clancey calls the frame of reference problem. It says that "theories are only true with respect to a frame of reference. AI and cognitive science research has been based on the contrary point of view that theories (representations and language) correspond to a reality that can be objectively known ..." [Clancey, in press, page 20]. There is no objective reality out there. By saying that there are no internal representations stored in any form, and that we can only know the things we can perceive, it follows that all knowledge is relative to the observer. Our perception is biased, and perceiving implies conceptualizing, interpreting. What we perceive is an interpretation in itself.

We agree that any science must operate within a certain frame of reference: the ontology it has chosen, the distinctions it deems relevant, the type of scientific justifications it allows and so on. However, this does not mean that scientific knowledge -and other knowledge for that matter- has no constancy or psychological validity. Both in KBS development and in education we strive to create systems or to coach pupils which *do tt right* according to some framework. Although, at the highest level, one may correctly say that all we know is relative to a frame of reference, within a particular frame of reference there exist agreed upon criteria of *right* and *wrong.* For instance, at school it is a particular framework of arithmetic and elementary mathematics that we are trying to convey to our pupils. Of course, the number facts and procedures are important, but the acceptance of the framework, i.e. learning that certain relations hold, certain solution methods are valid, for example accepting that 1+1=2 and not 3, is a major aim. Within that frame it is useful and necessary to talk about the adequacy of solutions being put forward by pupils, not solely in terms of right or wrong, but in terms of elegancy, efficiency, applicability or whatever criteria the frame of reference accepts. These criteria all are related to decontextualisation and naturally go together. The more abstract, formal procedures are the more elegant, more efficient, and more applicable ones. Thus modern physics teaching is aimed at changing fundamental preconceptions, e.g. that a force needs to be continuously applied to a mass in order for it to move with a constant speed, into the more abstract principles of physics. It is exactly this process which requires abstraction from the situation: decontextualisation. It is clear that human cognition is not very well equipped for the decontextualisation process: it is difficult, prone to error, and does not proceed autonomously. That does not mean, however, that one should not try to guide students through a process of decontextualisation.

Here a major flaw in the situationist argument appears to emerge. The fact that all knowledge is relative to a frame of reference, does not mean that humans have to be aware of that frame of reference, in fact most often they are not. Nor does the ever changing nature of frames of reference imply a parallelism in human memory which then no longer leaves room for stable symbolic structures. Working and thinking within a particular frame of reference provides stability and logical consequence to one's thinking. It is probably a good thing to teach students at some point in their career that these - different - frames of reference exist, but it is not at all a necessary condition for acquiring knowledge within a certain framework.

A major problem is the way in which we teach knowledge. How is knowledge applied?

5 Educational implications? How situated should our education be?

Brown et al. and Clancey claim that their revision of how to think of knowledge, representations and memory has strong implications for education. What do these implications boil down to? The central point is that knowledge is fluid, learning goes on all the time, learning entails creating new representations in speaking and acting and are part of every day life. Knowledge can never be simply transferred. People are actively creating representations while engaging in a community of practice. To learn a subject means to take part in a particular community of practice. This entails what is termed "cognitive apprenticeship". Educational researchers advocating this view on learning provide a lot of anecdotical evidence to support their claims [Resnick, 1988]; [Schoenfeld, 1985]; [Pea, in press]. What can we say about the promises of cognitive apprenticeship? What empirical findings do or do not support its claims? One of the consequences of the situated stance is to refrain from decontextualising, from abstracting the particularities of a problem situation. Formal education, as we know it, emphasizes decontextualisation. Its aim is to provide formalized tools that have broad applicability, their use is not restricted to a particular problem in a particular context. Let us take up an example provided by Brown et al 1989.

5.1 A situated example

In this example the inventiveness of a weight-watcher solving a fractions problem by other means than a formal procedure is described. The weight-watcher confronted with the problem of measuring out three quarters of two thirds of cottage cheese does so by making a round shape of cottage cheese, and subsequently dividing this shape into three parts, and then dividing two of these three parts in halves. Brown et al. claim this solution to be adequate, the situation the weight-watcher finds himself in provides him with a solution-path tuned to the problem at hand. There is no need to fall back on a formal solution. One could ask whether the solution described is indeed adequate. What if the weight-watchers material would have been some liquid? Then his procedure would break down. Formal multiplication of the fractions, however, would break down in neither case, and thus has a broader applicability [Elshout, in press]. Isn't that just what you would like to provide your students with: knowledge and procedures that are broadly applicable, decontextualised? Wouldn't cognitive apprenticeship have the risk that decontextualisation will never take place, so the student would become an able practitioner in a limited set of situations, without the possibility to solve problems in slightly different contexts? One could answer that formal education does not show much transfer of this kind either? That may be true. But, whether this lack of transfer is to be blamed on the efforts of decontextualisation is another matter. Could it be that lack of transfer is related to other factors than decontextualisation per se? Elshout [in press] argues convincingly that " at the

root of failure to transfer lies failure to comprehend complex information". The subjects taught at school are inherently difficult. The more abstracted the methods that can be used to solve problems, the more difficult the problems that can be solved. It is evident that many of us do not digest the material offered at school in full. Knowledge remains partial, half remembered and not completely understood. Cognitive apprenticeship in itself is no remedy. It does not reduce the inherent complexity of the subject. And it may be feared that it only conceals this complexity, by offerering the opportunity to "solve" a problem in a situated manner, like the weight-watcher, and thus failing to abstract from the particular context.

5.2 More on cognitive apprenticeship

Clancey is rather optimistic about the questions that students in cognitive apprenticeship would be able to pose [Clancey, 1990a]. Instead of simply taking in what is being put forward as the expert view, they would become aware of the differences of opinion that characterize the field. For a medical student this would mean asking questions like: if I know something, what other people will know it too, nurses or only particular doctors, or what if I were to travel to another part of the world, would I be able to talk to practitioners of medicine in the same terms? Rather high-brow questions for a student to ask? Is Clancey not expecting too much. Would a student in a cognitive apprenticeship not lack the necessary background experience to be able to pose such questions, and would he or she not be overwhelmed by the complexities of the field in action? Are there not just as many sound arguments against cognitive apprenticeship as there are in favour?

Formal education is not available world-wide. To many so called traditional societies, cognitive apprenticeship is the only educational path available. How does this compare to western, formal education? Is there any reason to claim that the cognitive apprenticeship that exists today is in any way superior to the formal education we provide our students with? We do not see one.

6 Conclusion

This paper argues that the situated stance in both AI and in education makes unfounded claims. Clancey arrives at the situationist stance by attacking the identity-hypothesis of which we say that it was never widely held anyway. The educators seek support for the situated view of education in case studies of cognitive apprenticeship. We find that these examples hardly justify a strong preference for situated learning. As we have said, decontextualisation does not come naturally to people. This supports the situated view, as it reflects how people normally act and learn in everyday life. However, it does not necessarily entail that formal education has to proceed in the same manner. On the contrary, formal education exists to free students from particular contexts. Without formal education, the necessary abstraction would not take place.

The conclusion of this paper must be that much of the situated cognition work today appeals to throw away the baby with the bath water. It is all very well to postulate new architectures and mechanisms for cognition, but disregarding existing, viable theories on the grounds that they do not hold for the full range of complex phenomena, is too drastic as long as new theories do not have a wider explanatory power. In fact, there is a danger of reductionism, i.e. reducing the mind to a simple organism interacting with its environment and producing complex behaviour through the application of simple behavioural rules. Psychological evidence disfavours this view. It takes a long time to acquire even basic skills, to become an expert requires hard work and a long training period.

It is not our opinion that situationism is entirely wrong. Psychological evidence about everyday thinking certainly supports the view that thinking is largely situated. What we argued against is the complete shift of paradigm that situationists claim. Disregarding evidence and achievements of Cognitive and Instructional Sciences, and AI, is in our view overstating the issues. There is as yet enough room within the traditional theories, to advance towards a full theory of intelligent behaviours.

From an AI point of view it is, of course, perfectly legitimate to investigate roads leading to new architectures, but stating that AI, and indeed educational science, have been built on faulty assumptions is too extreme a view for us to accept.

References

[Agre, 1988] P.E. Agre. *The Dynamic Structure of Everyday life.* PhD thesis, MIT, 1988.

[Anderson, 1978] J.R. Anderson. Arguments concerning representations for mental imagery. *Psychological Review,* 85:249-277, 1978.

[Anderson, 1979] J.R. Anderson. Further arguments concerning representations for mental imagery: A response to hayes-roth and pylyshyn. *Psychological Revttw,* 86:395-406, 1979.

[Bartlett, 1977] F.C. Bartlett. *Remembering: A Study in Experimental and Social Psychology.* Cambridge University Press, Cambridge, 1977,

[Bransford and Franks, 1971] J.D, Bransford and J.J. Franks. The abstraction of linguistic ideas. *Cognitive Psychology,* 2:331-350, 1971.

[Brooks, 1991] R.A. Brooks. Intelligence without representation. *Artificial Intelligence,* 47:139-159, 1991.

[Brown *et al,* 1989] J.S. Brown, A. Collins, and P. Duguid. Situated cognition and the culture of learning. *Educational Researcher,* 18:32-42, 1989.

[Chase and Ericsson, 1981] W.G. Chase and K.A. Ericsson. Skilled memory. In J.R. Anderson, editor, *Cognitive Skills and their Acquisition.* Lawrence Erlbaum Associates, Hillsdale, N.J., 1981.

[Clancey, 1990a] W.J. Clancey. Invited talk at the DELTA, 1990.

[Clancey, 1990b] W.J. Clancey. Why today's computers do not learn the way people do. Paper presented at the annual meeting of AERA, Boston, April, 16-20, 1990.

[Clancey, in press] W.J. Clancey. The frame of reference problem in ai. In K. VanLehn and A. Newell, editors, *Architectures for Intelligence: The Twenty-Second Carnegie Symposium on Cognition.* Lawrence Erlbaum Associates, Hillsdale, NJ, in-press,

[de Groot, 1966] A.D. de Groot. Perception and memory versus thought: some old ideas and recent findings. In B. Kleinmuntz, editor, *Problem solving: Research, method and theory,* pages 19-51. John Wiley and Sons, New York, 1966.

[Dennett, 1981] D.C. Dennett. Cognitive wheels: The frame problem of ai. In C. Hook way, editor, *Minds Machines, and Evolution: Philosophical Studies,* pages 129-151. Cambridge University Press, Cambridge, 1981.

[Elshout, in press] J.J. Elshout. Formal education versus everyday learning. In E. de Corte, M.Linn, H. Mandl, and L. VerscharTel, editors, *Computer-Based Learning Environments and Problem Solving.* Springer-Verlag, Berlin, in-press.

[Ericsson and Poison, 1988] K.A. Ericsson and P.G. Polson. A cognitive analysis of exceptional memory for restaurant orders. In M.T.H. Chi, R. Glaser, and M.J, Farr, editors, *The nature of expertise.* Lawrence Erlbaum Associates, Hillsdale, NJ, 1988.

[Harel and Papert, 1990] 1. Harel and S. Papert. Software design as a learning environment. *Interactive Learning Environments,* 1:1-32, 1990.

[Marr, 1977] D. Marr. Artificial intelligence; A personal view. *Artificial Intelligence,* 9:37-48, 1977.

[Mettes and Pilot, 1980] CT.C.W. Mettesand A. Pilot. *Over het oplossen van natuurwetenschappelijke problemen (On the solving of physics problems).* PhD thesis, University of Twente, 1980. Enschede: Educational Center.

[Newell, 1982] A. Newell. The knowledge level. *Artificial Intelligence,* pages 87-127, 1982.

[Pea, in press] R.D. Pea. Distributed intelligence and education. In D. Perkins and B.Simmons, editors, *Teaching for Understanding tn an Age of Technology.* in-press.

[Resnick, 1988] L. Resnick. Learning in school and out. *Educational Researcher,* 16:13-20, 1988.

[Rosenschein, 1985] S.J. Rosenschein. Formal theories of knowledge in ai and robotics. Technical Report 362, SRI, 1985.

[Schoenfeld, 1985] A.H. Schoenfeld. *Mathematical Problem Solving.* Academic Press, Orlando, FL, 1985.

[Simon, 1980] H.A. Simon, Problem solving and education. In D.T. Tuma and F. Reif, editors, *Problem Solving and Education,* pages 81-92. Lawrence Erlbaum Associates, Hillsdale, NJ, 1980.

[Staszewski, 1988] J.J. Staszewski. Problem solving and education. In R. Glaser M.T.H. Chi and M.J. Farr, editors, *The Nature of Expertise.* Lawrence Erlbaum Associates, Hillsdale, NJ, 1988.

[Suchman, 1987] L.A. Suchman. *Plans and situated actions.* Cambridge University Press, Cambridge, 1987.