TO BE INTELLIGENT OR NOT TO BE INTELLIGENT: IS THAT THE QUESTION?

Dick J. Bierman, University of Amsterdam,
Roetersstraat 15, 1018 WB Amsterdam, The Netherlands

Summary

The history and current situation of computers in education is discussed in terms of the development of psychological theories. 'Intelligent' Tutoring systems have been claimed to represent a significant improvement over traditional use of computers in education due to their added 'intelligence'. The amount and the form of intelligence in several educational media is analyzed. It is concluded that the amount of intelligence that goes into traditional systems probably has been underestimated. The problems with student models and learner models, which are the core components of an ITS, are briefly discussed. It is proposed to replace the quest for 'intelligence' by a quest for 'efficiency' when it comes to practical application of computerized education.

1 INTRODUCTION: A HISTORICAL PERSPECTIVE

1.1 Behaviourism

Even before the arrival of computers educational designers had constructed 'learning machines' that produced sequences of simple teaching transactions. Each interaction was initiated by the machine, displaying a question like 'apple = a) prune b) poule c) pomme'. The student was supposed to respond by pressing one out of several buttons. The feedback and the next interaction were generally a function of the last response (fig. 1).

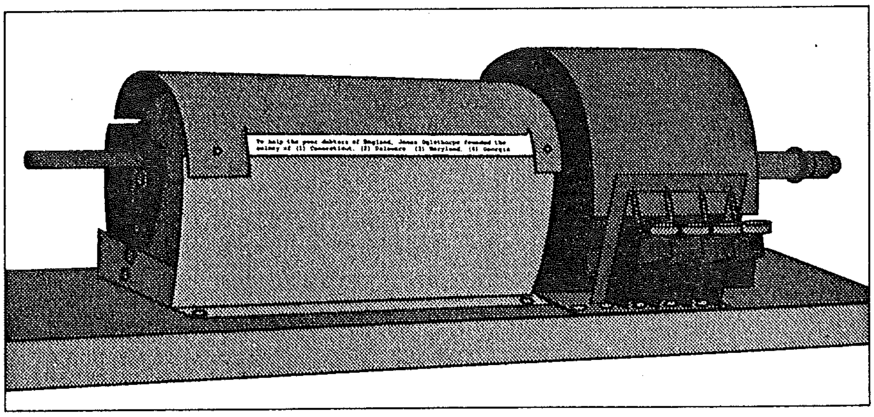

Pressey's 1927 device omitted questions previously answered correctly.

The machines were as much a reflection of then current technology as of the then prevailing psychological paradigm: Behaviourism. The whole history of artificial learning environments can be seen as an interplay between available technology and popular learning and teaching theory.

When the computer arrived on the scene, behaviourism with its stimulus-response paradigm was still the leading psychological theory. This theory explained how even pigeons could learn to produce simple responses, like picking on a button, when a stimulus such as a red light was presented. Complex behaviour could be trained by sequencing of simple responses. No wonder that the early educational computer software was structured similar to the old generation of non computerized learning machines. A lesson consisted again out of a series of interactions. For each unit of interaction, the stimulus, the possible responses, and the feedback for each of these responses was represented. Authoring systems that were used to develop courseware, usually offered the possibility to branch to other frames based on the history of student responses rather than on the recent response alone. Although this feature allowed authors to deviate from the basic behaviouristic premises that appeared to be valid for pigeons, they generally did not use this possibility. The vast majority of the programs from those early days can still best be described as simple drill & practice. Some authors did not even apply the limited and rigid interactivity of behaviouristic practice & drill. Their only model for teaching material was the book and so they mimicked books in hardly-interactive programs that were very appropriately called 'page-turners'.

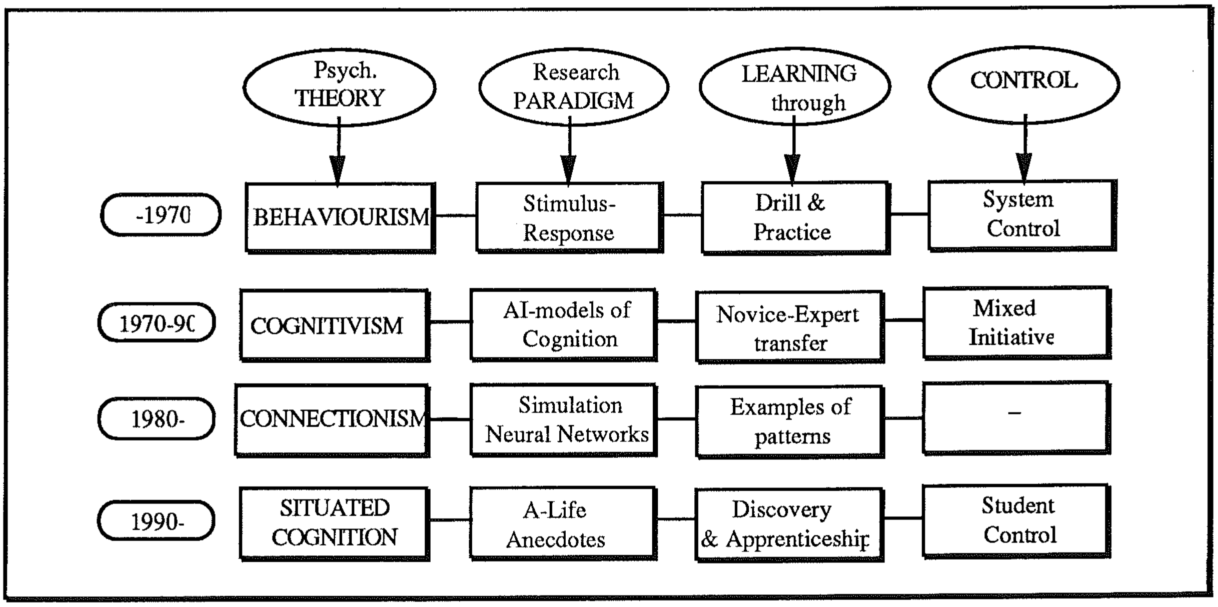

1.2 Cognitivism

The early period

In the 60's and 70's this situation changed gradually. Psychologists started to 'open' the black box in our skull. Behavior was no longer seen as a series of stimulus-reponse pairs but as a consequence of complex cognitive processes. Several psychologists formulated global cognitive mechanisms (Bruner, 1964, 1966; Ausubel, 1978), that is processes between the stimulus and the response, using knowledge as a basic intermediary concept. Based on these processes, principles of instructional design were formulated (Gagné, 1979; Merill, 1983; Scandura, 1983). These principles were applied in the design of courseware. Elements and structure of the course reflected the elements and structure of knowledge in the domain to be taught. Diagnosis would sometimes result in remedial teaching rather than in the behavioristic approach of presenting the same stuff again. Most of the CAI today is still designed with these two latter instructional principles in mind.

The AI period

The development of detailed cognitive theories was accelerated and stimulated by the developments in the field of Artificial Intelligence (and vice versa). The representation of knowledge as a distinct entity became focus of the research (Levesque, 1986). Introspection, which was totally denounced by the behaviourists, returned in the form of thinking aloud protocols. Subjects were asked for instance to solve a problem while thinking aloud and researchers tried to make sense of these data. The analysis of these protocols became a valid method to collect data on human knowledge representation and cognitive reasoning. When building models of cognition using AI-techniques the thinking aloud protocols served both as sources for the initial implementation of models as well as criteria for evaluation for the final implemented cognitive models. It appeared however to be difficult to formulate a formal method to evaluate AI-models of cognition in the way statistics became the formal method to evaluate quantitative data.

Part of the AI-community claimed --*that humans were nothing more than complex computers and that most intellectual tasks could be done by computers near the end of the century*--. Although it was never stated explicitly, this claim could be interpreted such that the task of teaching could be done as well by machines as by humans. Thus it was implicitly assumed by some that 'Intelligent' tutoring systems could replace the teacher in the classroom. To achieve the goals of building 'Intelligent' systems, like expert systems, much effort was put into the analysis of human experts knowledge representations and problem solving behaviour. Because knowledge now appeared to be 'stuff' that could be represented separately, education became the transfer of this stuff from the expert to the novice. So education was studied by examining the differences between novices and experts and the development in between. It should be stressed that not all ITS researchers shared the goal of building an artificial teacher. Many had the more modest claim that, by representing domain, diagnostic, and didactic knowledge explicitly the resulting systems would be more flexible than the traditional CAI systems.

The traditional, early cognitive and behaviorist, CAI programs were criticized eg. because generally these rigid systems did base their *pre-programmed* interventions on the last response only. In 'Intelligent' systems a dynamic model of the student was maintained and behavior of the system was *generated* using conditions from this cognitive model.

1.3 Parallel developments

Discovery Environments

Parallel to the AI triggered development of ITS systems, a separate reaction to the early cognitive CAI became apparent. These old CAI programs resulted essentially in systems initiated actions because teaching was seen as the transfer of knowledge from the system to the student. In 1961 already Bruner pointed to the possible cognitive relevance of the act of discovery (Bruner, 1961). Papert took up this lead and explicitly proposed that students would profit most from discovering themselves the basic principles of the domain (eg. mathematics or physics; Papert, 1980). The implicit pedagogical assumption was that, given the right discovery environment, in this case the LOGO environment, students on their own initiative and without any support would find their way. Thus a cognitive diagnosis would be unnecessary and no student model would have to be maintained by the system. Most research in this area used anecdotal evidence and it was very hard to evaluate the claims that discovery resulted in the acquisition of 'powerful (context independent) ideas'. In practice many discovery environments included some guidance of the student (Abelson & DiSessa, 1980).

Tools

Shortly after the LOGO-developments, the software technology provided users with several tools like powerful mathematical and statistical packages, simulation environments etc. In a few years, most computers in the classroom, especially in engineering and science, were used as *tools* rather than as artificial teachers (van Dijk, 198x). Not only the availability of these tools was instrumental in bringing about this shift in educational use of computers. This trend also was reinforced because early CAI programs (drill & practice / page-turners) tended to be very dull and programs with a more cognitive approach tended to be very expensive in terms of development.

1.4 Connectionism

By the end of the eighties the cognitivist paradigm was heavily questioned by what was labeled as Connectionism. Connectionists focussed on massive parallel distributed computations and simulations of neural networks. This paradigm succeeded to build systems that could learn to discriminate between patterns. The representation of knowledge in connectionist systems is by a set of weights while in AI-models of cognition it is by logical statements. Learning is reflected in changes of the weights and the sole learning mechanism is learning from examples of patterns.

Although Connectionism appears to deal well with primary and immediate cognitive processes, like recognition, it is difficult to see how connectionist systems can be used to model complex problem solving where students might get stuck and restart or use 'repair' strategies.

1.5 Situated Cognition

Whether it was the appearance of the new fad 'connectionism' or the disappointment with the obvious failure to build really intelligent systems before the end of the century is unclear, but many AI-based cognitive scientists turned away from their roots. Most notably the 'father' of educational expert systems, W.C. Clancey (Clancey, 1990?) and the 'father' of student-model (malrule-) based ITS, J.S. Brown (Brown, 1990?). They currently support the approach advocated by the field of Situated Cognition. Situated Cognition denounces the basic AI premise that knowledge can be seen as some context independent 'stuff' that can be 'poured' in a student. On the contrary, according to the situated cognitive scientists, all knowledge only becomes knowledge in some context and it is impossible to represent this in a subject's head. The field of situated cognition is so recent that it is difficult to infer the implications for computers in educations. It appears however that some older ideas in education like project-education, on the job training, and cognitive apprenticeship re-emerge. The role of the computers could become more one of a communication tool linking learners to eachother and to real contexts. Anyway, according to situated cognition theories, most education should be in real contexts like the workplace and the knowledge transferred should be very contextualized eg. "expertise stories" (Orr, 1986; Bransford, 1987). Less outspoken views hold that this might be a fruitful approach for a few domains (eg. motor- and other automated skills like skiing) or parts of domains while in other areas the idea of transferring knowledge to the student is still a valid approach (Sandberg & Wielinga, 1991).

2. WHERE IS THE INTELLIGENCE?

2.1 Non-interactive environments

An educational environment consists of a set of materials and humans which can interact to achieve certain educational goals. In the following we will analyze how much 'intelligenge' goes into such an environment.

Let's consider a textbook. A textbook is not a formal description of the domain. A formal description would in itself suffice as a complete description and could be constructed on the basis of an analysis of the domain alone. A textbook is more. The difference is clear for the domain of programming languages. Most programming languages come with a (rather formal) reference manual. Nevertheless there is obviously a need for another representation when it comes to teaching this programming language. In the instructional textbook the author has not only the domain but also the student in mind. The material is presented so that learning will be optimal. Since the textbook is the same for all students the author has an average student in mind. In fact a textbook is the frozen representation of lots of intellectual labor of the author in which a non-specific model of the student was used. For instance the author will estimate the average knowledge-level of the student at the beginning of the book. The intelligence of the author of a textbook also might be used to anticipate upon common misconceptions and try to provide counter-examples. The important ongoing question for each educational system is *what and how to present it next.* It is in answering this question that the intelligence is needed.

Of course students without the anticipated misconception will go through a counter-example which is more or less unnecessary for these particular students. Or students that have capabilities above the assumed average student might become bored. In other words, the decision about *what to do next and how to do it* was for these students sub-optimal.

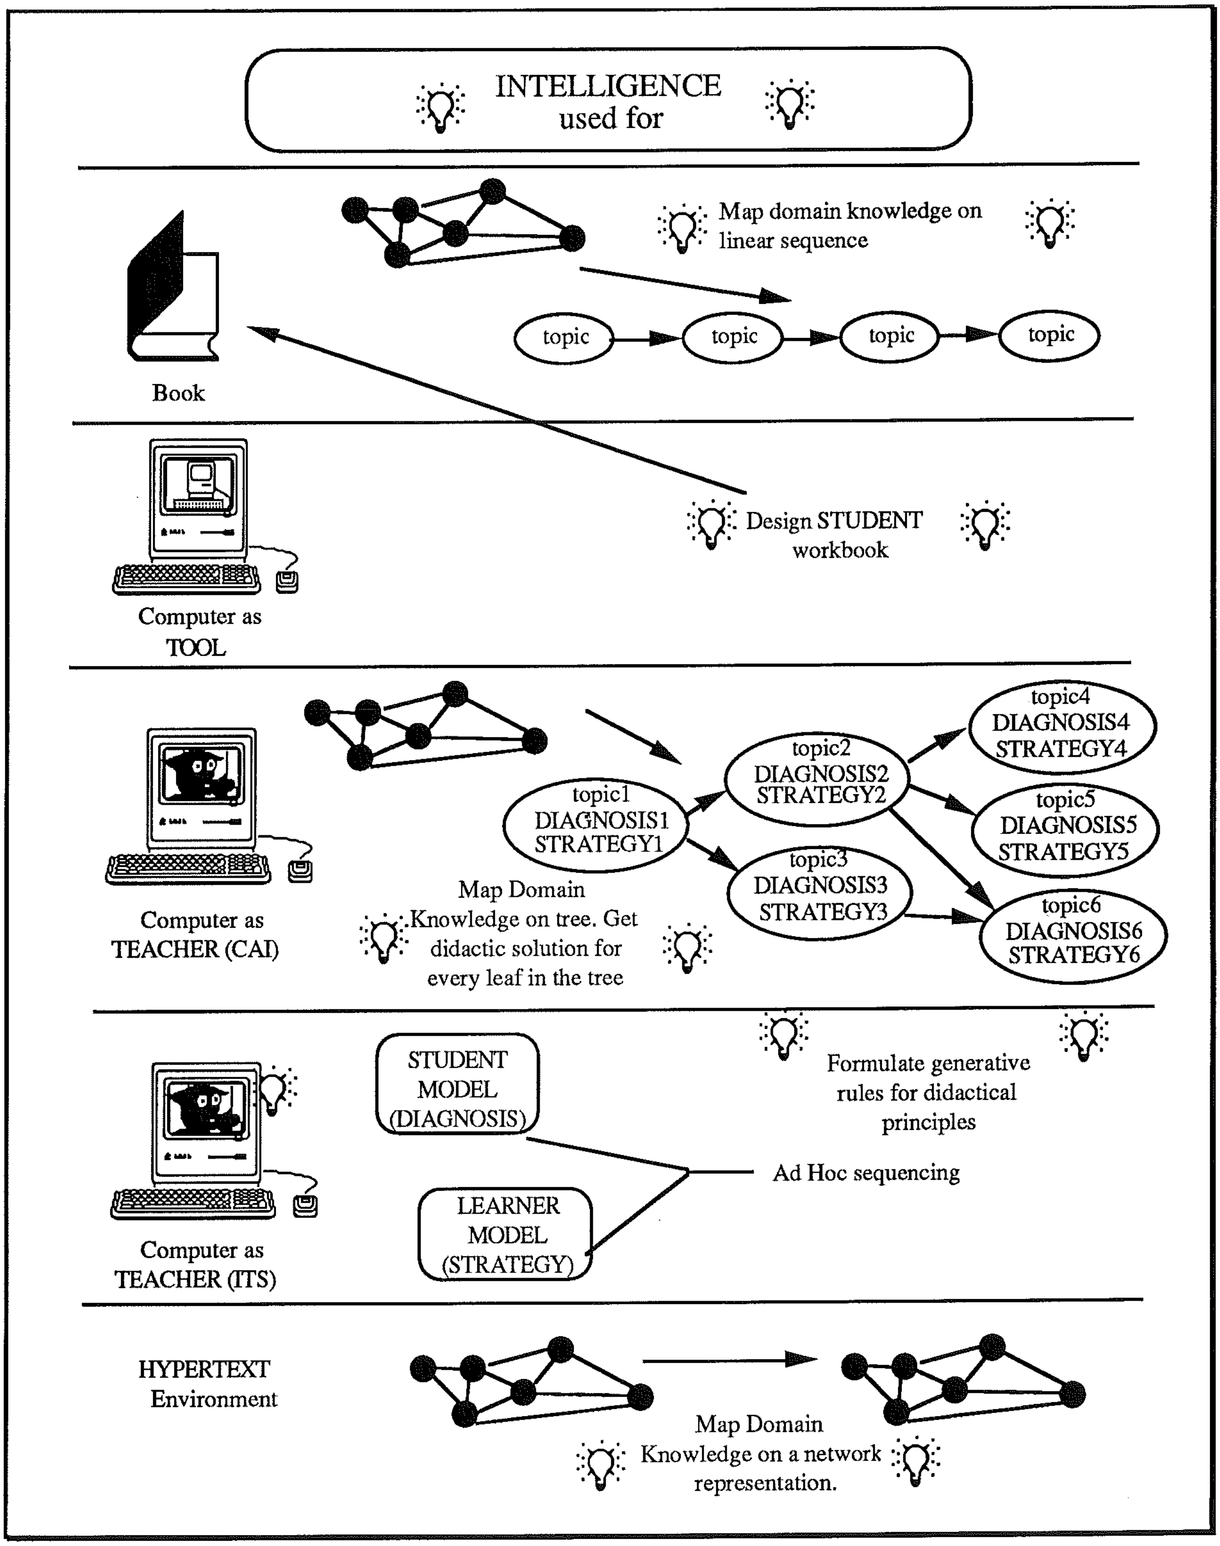

2.2 Interactive environments

But if we come to traditional CAI the superfluous counter-example can be avoided. A good author of traditional CAI does almost the same as the author of a textbook but instead of anticipating an average student, several types of students are anticipated. The work a courseware-author puts into course design can be compared to the labor of a chess player studying his opening repertoire. Trying to anticipate each move of the opponent, to analyze its meaning (plan), and to find the best response. What these people have in common is that they try to solve as much of their problem beforehand. The diagnostic knowledge, and the knowledge of teaching strategies that the author uses in solving this problem is not explicitly represented, only the solutions are. The knowledge therefore is embedded in the concrete contexts of the resulting series of teaching events (fig.3). Of course the constructed artifact (opening repertoire or CAI-program) in itself is not intelligent. The intention here is to show that they both are the result of intellectual efforts and hence their performance might be comparable to artifacts like ITS that are labeled 'intelligent' because they can generate (very often the very same) solutions.

In terms of the chess player, traditional CAI knowledge is restricted to the opening repertoire, it can only deal with situations that have been predicted. 'Intelligent' tutoring systems have generative knowledge. An ITS can dynamically generate the solution of the question which educational action is the best in the didactic circumstance even if the situation has not been predicted in detail.

2.3 Discovery environments

Most intelligence, either used by the author in preparation of, or embedded in, a system results in better decisions on the 'what and how' of subsequent teaching actions. It is implicitly assumed that control about the 'what and how' can be executed by the system. For discovery environments however, it is not the system but the student who is in control. This also holds for the situation where the computer is used as a tool. In these cases the intelligence generally does not reside in the computer but elsewhere within the educational context. For instance in sessions where the computer is used as a tool, students usually have some form of a workbook with a clear suggested sequencing of the problems to be solved. Often there is also some assistant walking around. Current research projects like HYPOCAMPE (Forte & Bierman, 1991) try to add some intelligence to the 'computer as a tool' environment by diagnosing the students behavior. From that diagnosis the entry point and recommended path in a hypertext environment is determined. The implementation of the hypertext is such that students have the idea they are in control but in reality there is hidden control by the computer.

2.4 Collaborative environments

Recent research on collaborative CAI has focused on optimalisation of the composition of the dyads of students by use of cognitive parameters. The idea was that part of the needed intelligence could be drawn from the co-student. Preliminary results, derived from a too small sample size, indicate that students behaved in accordance with the prediction that the (introvert) more able students would be forced into the role of teacher. However, contrary to what would be expected, the data suggest that this behavior did not result in better learning for the less able student (Balk & Bierman, 1989).

3. STUDENT MODELS AND LEARNER MODELS

How is it possible to know what teaching action, or treatment, in a given context is best? The diagnosis and representation of the student's cognitive state is the kernel aspect of most ITS. In most architectures this representation is the database where all components refer to. At any moment 't' in the learning history this cognitive state is called the (dynamic) 'Student Model'. The problem of how to select the appropriate optimal teaching action or treatment, given that the Student Model at that point in time is known, has been given less attention. It was realized that to solve this latter problem a detailed theory on individual learning should be available. Such a theory should then be implemented as a Learner Model. A sharp distinction between Student- and Learner Model is necessary: The Student Model is a series of snapshots of the students' cognitive states. The Learner Model describes how individuals learn. Recently it has been proposed that given the Student Model at time 't' and a model of how learning proceeds the system could calculate the Student Model at time 't+1' for all different possible treatments (VanLehn, 1991). Then the system should be able to decide which of the possible treatments was the best.

3.1 Practical and fundamental problems

The natural temporal order, where diagnosis precedes treatment, is not the only reason that explains why research has focussed on Student Models. The fact that there are little or no good theories of individual learning that can be implemented as a model of the learner is another one. Instead of the generative approach based upon running a Learner Model, as described in the previous paragraph, the choice of treatment also might be done directly on the basis of the Student Model. In that case expert didactic knowledge derived from human teachers is used to infer the subsequent treatment. The assumption is that good human teachers implicitly 'know' how students learn. Since this knowledge is 'compiled' they can not explicitly formulate this in the form of a Learner Model. But rules might be deduced from the actual behavior in specific didactic situation which reflect their Learner Model. It has been tried to elicit these rule by analysis of thinking aloud protocols of teachers in one-to-one tutoring session but the results are very meager (Kamsteeg & Bierman, 1991) .

• *Student models: the bandwidth problem.*

Soon after it became clear that diagnosis of a student was crucial in ITS design, it was realized that there was a fundamental constraint on this process of cogntivie diagnosis. A Student Model could never be assessed with enough detail due to the limited bandwidth in the communication between the student and the system (Goldstein, 1982). Some research has focussed on the construction of interfaces which increase this bandwidth and permit the observation of more intermediate problem solving steps (Bierman, Kamsteeg & Sandberg, 1991). Part of this problem also might be alleviated by having the learner model calculate those intermittent cognitive states of the student for which the system lacks direct information. Of course it is not possible to calculate the Student Model from scratch since the initial state at t=0 is unknown. So initial cognitive diagnosis remains necessary.

• *Learner models: problematic fit on studentmodels; poor treatments*

Most models of learning in the field of computer learning are considered to be psychologically invalid. And even if these models could be used as a Learner Model in an ITS it is not a straightforward task to map their representations of the domain knowledge onto the representation of knowledge in Student Models. The treatments, to train these systems, consists of examples and/or counter-examples. In real education there is a richer variety in treatments. It is difficult to formalize these rich treatments especially if natural language plays an important role.

SOAR, which is a more recently developed model for learning, claims psychological validity (Simon, 1991). For this particular Learner Model the problem of mapping knowledge representations is prominent. Knowledge in SOAR is 'hairy' and difficult to inspect. Furthermore it is computationally impossible to run SOAR for each possible treatment and still have reasonable response times.

3.2 Empirical data

3.2.1 The extra value of dynamic student models

If the dynamic Student Model is the basic element in ITS then it should follow that systems where this element is replaced by a static student model are less efficient when it comes to teaching. It is astonishing that hardly any relevant empirical research exists that tries to evaluate the extra value of a dynamic student model. Most experiments do not compare the dynamic model with a static model but with no model at all. In other words, even the intelligence that is reflected in textbooks is removed. Furthermore they do not control for time on task. Sleeman did an experiment where he compared treatments based upon a cognitive diagnosis with a uniform treatment of showing the correct solution (Sleeman, 1989). No difference between the two treatments could be asessed. This experiment has informally been criticized because of the choice of an improper domain (vanLehn, 1991) and because of the use of improper statistics (Elsom-Cook, 1991). These objections are rather ad-hoc. The only reasonable objection seems that the cognitive treatment is rather crude. It boils down to just comunicating the cognitive diagnosis to the student. It is the psychoanalytical assumption that one will 'heal' as soon as one knows the problem. This assumption is highly questionable but on the other hand there are no other obvious treatments available.

3.2.2 Non verbal and Affective components?

There is ample evidence that computer assisted education is more effective than classroom teaching. Its results in terms of efficiency might be compared with human teaching of small groups. However, when comparing computer assisted teaching with individual human coaching it apppears that the human is more effective than the computer, although the evidence is not overwhelming (Bloom, 1984). From experiments where protocols were taken from human coaches, it might be concluded that these human teachers did not maintain an elaborate and dynamic student model (Kamsteeg & Bierman, 1991). Most 'what and how to do next' decisions of the human coach were based upon the last action of the student and a general classification (like 'weak student'). There might be three, not necessarely mutual exclusive, explanations of the finding that human coaches still outperform computers as teachers.

In the first place the human coach is not bothered by the restricted bandwidth. The coach observes several non-verbal cues indicating emotional and cognitive states like the student being bored, angry, completely lost etc. This gives the human coach much extra information on when to intervene.

Secondly the coach herself might display a few affective states that might have impact on the student. Most notably 'enthousiasm'.

Thirdly, human coaches might us more effective instructional methods, presenting a remedy for a multitude of possible learning problems in one transaction (Luarillard, 1990)

4. CONCLUSION

Although not much empirical evidence is available, it appears that the 'intelligence' of ITS does not add much in terms of performance of courseware. This is consistent with the analysis of the amount of 'intelligence' that goes into the design and implementation of traditional good educational material.

It is therefore highly questionable if the extra work that is needed to build these ITS is worth the effort. It might have been if this work could be re-used. It has has been suggested that due to the separation and explication of the different types of knowledge an ITS for domain X could be constructed from an ITS for domain Y by simply replacing the domain knowledge. This turns out to be a far too optimistic suggestion due to the non-orthogonality of the didactic- and domain knowledge. When it comes to practical applications the conclusion seems justified not to invest too much in the ability of the system to generate optimal teaching transactions on the basis of a student and/or learner model. Instead of striving for maximal intelligence it is recommended to strive to maximal efficiency.

Some of the problems with ITS sketched in the preceding paragraphs might be fundamental. This does not necessarily mean that ITS research is at a dead end. There is already much spinoff. Architectures and ideas from ITS are penetrating the field of traditional CAI (Balk & Bierman, 1990; Sandberg & Barnard, 198x). Research on dialogue systems is triggered by

ITS-research. And recently a Learner Model was used to evaluate traditional media like books (vanLehn, 198x). The Learner Model of SOAR was used off-line to evaluate the user friendliness, in terms of learning, of different computer interfaces (Young, 1991). These evaluations might result in improvements of the materials.
ITS appears still the royal avenue to more understanding about individual learning because it offers a way to explore and evaluate Learner Models. In doing so ITS research also might have impact on new developments in AI, more particulary in machine learning. Research could and should proceed even if we know that it will never result into an ITS with a Student Model and a Learner Model which will actually work in the classroom.

Acknowledgement
Paul Kamsteeg has contributed a significant part of this paper by pointing out that the history of psychological theories and its implications for computers in education is more complex than I originally thought it was.

REFERENCES

Abelson, H., and DiSessa, A.A., (1980). Turtle geometry, the computer as a medium for exploring mathematics. Cambridge (Mass): MIT Press.
Anderson, J.R., Conrad, F.G. & Corbett, A.T (1989) Skill acquisition and the LISP Tutor. Cognitive Science, **13**, 467-505.
Ausubel, D.P., Novak, D.J., and Hanesian, H., (1978). Educational psychology, a cognitive view (2nd ed). Cambridge (Mass.): MIT Press.
Barnard, Y. & Sandberg, J.A.C. (1988). Applying artificial intelligence insights in a CAI program for 'open sentence' mathematical problems in primary schools. Instructional Science, 17, 263-276.
Bierman, D.J. & Balk, V. (1989) *Towards a cognitive theory of Student Cooperation in CAI.* Proc. of the 31 st Int. Conf. of the Association for the development of Computer-Based Instructional Systems, pp.220-222. ADCIS, Western Washington University, Bellingham, WA.
Bierman, D.J. & Balk, V. (1990). A general purpose architecture for traditional CAI using an advanced authoring system. Proceedings APPLICA '90, p.1-8, Lille, sept. 90.
Bierman, D.J., Kamsteeg, P.A. & Sandberg, J.A.C. (1991) Studentmodels, scratch-pads and simulation. Proc. of NATO ARW, 6-10 oct. 1990, Sintra, Portgal.
Bloom, B.S., (1984). The 2 sigma problem: the search for methods of group instruction as effective as one-to-one tutoring. Educational Researcher, 13, 3-16.
Bransford, J.D. (1987). Enhancing thinking and learning. New York: Freeman.
Brown, J.S. , Collins, A.,& Duguid, P. (1989). Situated cognition and the culture of learning. Educational Researcher, 18, 32-42.
Bruner, J.S., (1961). The act of discovery. Harvard Educational Review, 31, 21-31. Bruner, J.S., (1964). The course of cognitive growth. American Psychologist, 19, 1-15.
Bruner, J.S., (1966). Toward a theory of instruction. Cambridge (Mass.),: Harvard University Press.
Collins, A., Brown, J.S., & Newman, S.E., (1989).Cognitive apprenticeship: Teaching the craft of reading, writing,and mathematics.In: L.B. Resnick (Ed.), Knowing, learning, and instruction: essays in honor of Robert Glaser. Hillsdale (NJ): Erlbaum.
Clancey , W.J. (1990). Why today's computers do not learn the way people do. Paper presented at the annual meeting of AERA, Boston, April 16-20, 1990.
Dijk, T van., Gastkemper, F., Moonen, J., & Romeijn, W., (1985).Motives for CAI in post-secondary education.Journal of Computer-Based Instruction, 12, 8-11.
Elsom-Cook, M. (1991). Discussion at NATO ARW on student modelling. Ste. Adèle, Quebec, Canada from May 5 to 8, 1991.
Forte, E. & Bierman, D.J. (1991).Implementation of student-sensitive Hypertexts as Help-resource in Test- and Simu-ware. Unpublished Manuscript.

Gagné, R.M., & Briggs, L.J., (1979). Principles of instructional design (2nd ed.). New York: Holt, Rinehart & Winston.

Glaser, R., (1984). Education and thinking: the role of knowledge. American Psychologist, 39, 93-104.

Goldstein,I.P. (1982). The Genetic Graph: a representation for the evolution of procedural knowledge. In D. Sleeman & J.S. Brown (Eds). Intelligent Tutoring systems. Academic Press, New York, pp.69.

Kamsteeg, P.A., & Bierman, D.J., (1991). Teacher thinking, student thinking: Think aloud protocol analysis in tutoring systems research. To be published in: P.Goodyear (Ed.), Teaching Knowledge and Intelligent Tutoring, (ch.9). Norwood, Ablex. In Press.

Laurillard, D., (1990). Generative student models: The limits of diagnosis and remediation. In: M. Elsom-Cook, (Ed.), Guided discovery tutoring: A framework for ICAI research. London: Chapman.

Levesque, H.J., (1986). Knowledge representation and reasoning. Annual Review of Computer Science, 255-287.

Merrill, M.D., (1983). Component display theory. In: C.M. Reigeluth (Ed.), Instructional-design theories and models: An overview of their current status. Hillsdale: Erlbaum.

Orr, J. (1986). Talking about machines: social aspects of expertise. Proc. Conf. on Computer Supported Cooperative Work. Austin.

Papert, S., (1980). Mindstorms: Children, computers and powerful ideas. New York: Basic Books.

Scandura, J.M., Instructional strategies based on the structural learning theory. In: C.M. Reigeluth (Ed.), Instructional-design theories and models: An overview of their current status. Hillsdale: Erlbaum.

Sleeman,D. , Kelly, A.E., Martinak, R., Ward, R.D., Moore, J.L. (1989). Studies of diagnosis and remediation with High School Algebra Students. Cognitive Science, 13, 551-568.

VanLehn, K. (1991).A Pseudo-Student that Learns Physics and its Implications for Instructional Design. Presented at NATO ARW on student modelling. Ste. Adèle, Quebec, Canada from May 5 to 8, 1991.