Intelligent Tutoring Systems: Learning and Cognition

Course Paper "Cognition, Learning, and IT" (Part 2)

Dennis Maciuszek

April 5, 2002

1 Introduction

Practical evaluation In [Maciuszek, 2002] I evaluated the e-learning software ALEKS[1] according to cognitive criteria. ALEKS was classified as an *intelligent tutoring system (ITS).* The intelligence is in its automatic adaptation to the learner's previous and learned knowledge. Thanks to successfully implemented adaptivity, the virtual tutor can provide individual support during the process of learning an area of maths. A comprehensible interface and a clever feedback concept contribute to a learning situation resembling school, but with a personal coach. However, there is almost no relation between the taught material and real world practice. All you do is solve abstract mathematical problems. The tutor's support of only a "bottom-up" learning style was further critisised. You learn to solve little problems, but are not told what their relevance for the whole area is. Neither is such holist knowledge being adapted to.

Theoretical analysis This essay seeks to complement the practical evaluation of ALEKS with a theoretical discussion of the ITS concept as one type of e-learning. From the perspectives of modern theories of learning and cognition, what is so good about intelligent tutoring? What was found problematic in the practical evaluation, and what could ITS do better? I will use "ITS" as a collective term for e-learning environments automatically adapting to knowledge, but limit myself to analysing mostly the approach ALEKS is built on, namely Doignon and Falmagne's theory of *knowledge*

[1] `http://www.aleks.com`

spaces [Doignon and Falmagne, 1999]. ITS in that tradition assess a student's knowledge by letting her or him solve a small number of exercises (*knowledge diagnosis*), and then build a student model (*knowledge state*) of the overall exercises (*items* or *problems*) in the domain (represented by the *knowledge space* of all possible knowledge states) which the student can solve. The student model allows the presentation of appropriate exercises at the student's precise level of knowledge.

2 What ITS Can Do

2.1 The Cognitive Perspective

The cognitve revolution Influenced by the ideas of rationalism and the appearance of Computer Science and its models of information processing, Cognitive Psychology originated in mid twentieth century to successfully challenge contemporary behaviourism (cf [Säljö, 2000], Chapter 3). The major achievement of the so-called cognitive revolution was the discovery of complex processes in the human mind. Instead of simple "stimulus ⇒ response" schemes – which undoubtedly account for primitive parts of human learning (*reactive* in the terminology of [Norman, 1993], *procedural* in that of [Tulving, 1985]) – there were now "stimulus ⇒ mediation ⇒ response" schemes (referring to [Cobb and Bowers, 1999], page 4). The "mediation" part would stand for all sorts of complex, cognitive processes.

Cognitive modelling Ohlsson in [Ohlsson, 1995] models human problem solving, like finding routes or playing chess, by *rule based* traversing of *situation trees*. He emphasises though that such *practical knowledge* is not the only kind. Ohlsson is looking for logic models to explain *declarative knowledge*, such as interpreting a map or knowing the laws of motion.

The ITS perspective The cognitive perspective on knowledge is decidedly analytical. Cognitivists model the human mind and domains of knowledge by detailed mathematical structures in order to approach their complexity. Learning is thus seen as acquisition and re-arrangement of an individual's knowledge structures (or beliefs, see [Ohlsson, 1995]). It is obvious that the cognitive perspective of learning has inspired the work on ITS. Intelligent tutoring intends to provide teaching tailored to the individual student, and it is common that the software builds an internal student model. ITS based on knowledge space theory work with mathematical models of student knowledge that are derived from the mathematical model of the respective domain.

Figure 1 shows this relationship, thereby abstracting from too much mathematical detail. For the scope of this essay, I am working with this figure as a simplified meta-model of a generic ITS.

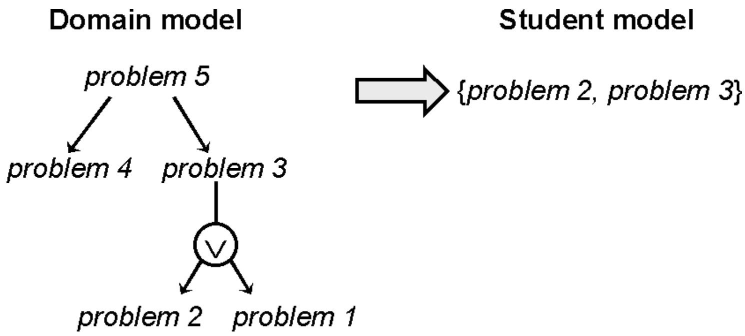

Figure 1: A generic ITS (1)

Modelling the student Cognitive theories regard individual learners as *agents* ([Greeno, 1997], page 7), who "perform tasks" ([Ohlsson, 1995], page 41). In knowledge space theory, an individual learner is represented by the set of domain problems (ie tasks) he or she is capable of solving (or performing).

Example 2.1 (Student knowledge) *Assume a very small maths domain "adding fractions greater than 1". One learner, somewhere at an intermediate level in mastering the domain, may be capable of writing* fractions > 1 *and writing* mixed numbers *like* $1\frac{1}{2}$.

Modelling the domain Instruction design inspired by cognitive learning theories, usually decomposes domain knowledge into smaller units or tasks. [Greeno, 1997] refers to this structuring of detailed domain knowledge as "factoring" (page 9), [Anderson et al., 1997] as "task analysis" (page 21). An ITS need not have a domain model, but this idea is central to knowledge space theory. The designer of an e-learning course – with the help of experts, experimental data, educational texts, or deeper analyses – divides domain knowledge into a set of items or problems of various difficulty. Prerequisite relationships are established between the items, so that it follows which sets of items make allowed knowledge states, ie consistent student models. The set of possible knowledge states is the knowledge space. Learning means acquisition of new items, thus reaching new knowledge states.

Example 2.2 (Domain decomposition) *The small domain "adding fractions greater than 1" may be decomposed into five items* add fractions > 1 *(the most difficult problem),* add fractions < 1, fractions > 1, mixed numbers, *and* uncommon fractions $(\frac{3}{2})$. *Prerequisite relationships can then be established, so that the knowledge space follows. Figure 1 shows prerequisite relationships as implication arrows. If a learner can solve* problem 5*, it follows that he or she can solve* problem 4. *Hence,* problem 4 *is a prerequisite of* problem 5. *In terms of this example: If a learner can solve* add fractions > 1*, it follows that he or she can solve* add fractions < 1.

Benefits The cognitive approach of ITS introduces the artificial intelligence of "virtual teachers". Student modelling allows a piece of software to "get to know" its user and to tailor teaching to her or his individual needs.

I am convinced that student modelling should be done in relation to a domain model, eg as described above. Every domain is different, not just in its concepts, but also in its cognitive nature. [Norman, 1993] and [Tulving, 1985] agree that there are at least two kinds of knowledge. *Experiential/reactive* (Norman) or *procedural* (Tulving) knowledge is automatic and not conscious. Higher knowledge is *reflective* (Norman). Tulving differentiates further between facts about the world, which are stored in the *semantic* memory system, and knowledge related to personal experiences, which is stored in the *episodic* memory system. Tulving agrees with Norman's three kinds of learning, namely *tuning*, *restructuring*, and *accretion*. Dreyfus and Dreyfus in [Dreyfus and Dreyfus, 1986] however relate kinds of knowledge to skill levels rather than domains.

In addition to gaining the structure of domain knowledge, task analyses immediately provide problems (cf Figure 1) to build interactive exercises on. ITS like ALEKS use such exercises both for the diagnosis of known items – which need not take long due to inferences based on the prerequisite relationships – and for teaching. Learning with ALEKS is "learning-by-doing", allowing immediate practise of learned material – an effective technique according to [Anderson et al., 1996]. The more precise the tasks are defined, the more differentiated feedback the software can grant to student's answers in exercises. Feedback from a coach provides a learner with essential reflection on her or his performance (see [Norman, 1993]).

Limitations The cognitive analytical view of learning has natural limitations. Decomposition of learning domains will never be without loss. Firstly, no mathematical model is able of adequately represent the complexity of the human mind. Nevertheless, some processes and states can be described in a pragmatic fashion. Cognitivist theories do centre around the individual human learner, which is hardly a bad cause as argued in [Anderson et al., 1997]. Secondly, no mathematical model can ever describe the complexity of the real world. Even the most advanced problems in ALEKS remain problems of "school culture" (cf [Seely-Brown et al., 1989], page 34). I would claim that this need not be, and that improved ITS models could describe more complex situations.

I support the cognitive view of learning in so far that I believe in detailed mathematical models that can pragmatically simulate learners and knowledge for ITS, at least attempt to clarify complex learning situations (cf [Anderson et al., 1997]), and provide a feeling for what is beyond the model – ie what is lost.

2.2 Adaptation to Knowledge

Requirement Many pieces of educational software aim at groups of users with a certain level of previous knowledge. A typical audience are beginners. Advanced learners using such software would soon get bored by too easy lessons. Likewise, beginners experience frustration, if a course is too difficult. Considering [Dreyfus and Dreyfus, 1986], such frustration reaches deeper than the lack of previous knowledge. Dreyfus and Dreyfus identify five levels of knowledge, and thus five types of learners, from novice to expert. These differ not only in the amount of domain concepts known, but also in the kind of representation, usage, and acquisition of knowledge. In short: Novices (level 1) learn facts. Advanced beginners (level 2) become sensitive to context. Competent learners (level 3) reason. Proficient learners (level 4) become personally involved. Experts (level 5) perform difficult tasks by grounded intuition. The authors themselves denote the question whether artificial intelligence can really represent higher levels of knowledge.

Idea The idea of ITS automatically adapting to the user's knowledge is to provide learning support for many different levels. Knowledge space models do not know the discrete five levels Dreyfus and Dreyfus propose. Their transitions can be smooth, due to detailed task analysis. It is up to the designers to judge, if they can identify problems from Dreyfus and Dreyfus' level 1 to level 5, and then create appropriate lessons and exercises. In many cases, student models covering the smooth transitions between two or three levels, eg from fact knowledge to reasoning, should suffice for the assessment of many students' knowledge states.

After the model of student knowledge has been built, eg by an initial diagnosis, the ITS is able to present lessons and exercises of *subjectively* average difficulty. The software can *prescribe* or *recommended* ([Ford, 2000], page 554) the most appropriate next lessons. ALEKS' strategy is somewhere in the middle of those two, as it allows the user to choose between possible prescribed next steps. The practical evaluation found the prescriptions consistent with the domain structure and student model, indicating that automatic adaptation to knowledge works.

Benefits The direct conceptual advantage of adaptive learning environments in comparison with conventional e-learning software is their ability to satisfy the learning needs of the individual student. The software "knows" precisely at what level to teach.

Moreover, I see positive implications beyond the intended. The design of lessons and exercises – the content of the problems – can hugely benefit from knowing the domain model and possible knowledge states. In a system that prescribes paths through the knowledge space, the prerequisites of an exercise being taken can be considered known by the student. Lessons could support the meaningful integration of the new material in the existing knowledge. [Craik and Lockhart, 1972] would call this a deep *level of processing*, resulting in better rememberance later. The already known item would as a side effect be *elaboratively rehearsed*, ie processed on a deeper semantic level than before.

[Norman, 1993] stresses the importance of intrinsic motivation in learning – by the name of Csikszentmihalyi's *flow*. Learning may lead to joyful experiences just by doing it, by taking part in the process. Norman lists seven factors that contribute to rewarding experiences of flow. Many of these, like intensive feedback, specific goals, tools that fit the user, are characteristics attributed to ITS exercises. In particular, I see the well-designed ITS

> "provide a continual feeling of challenge, one that is neither so difficult as to create a sense of hopelessness and frustration nor so easy as to produce boredom."
>
> [Norman, 1993]

Limitations Besides its cultural limitation to the school context, I critisised the "bottom-up" learning style demanded by ALEKS. It may suit some learners, yet discourage others. Knowledge is not the only inter-individual difference in learners. *Cognitive styles* should be added. Normally, these are not supported by student modelling based on knowledge space theory.

The "bottom-up" learning style is discussed (under different labels) by both [Sternberg and Grigorenko, 1997] and [Ford, 2000]. Starting with little details as the smallest resulted pieces from task analysis seems to be inherent to the idea of learning with ITS adapting to knowledge. Again, I wonder if this is necessarily the case. In the next chapter, I make suggestions for the better design of ITS and possible extensions to the concept presented so far.

3 What ITS Could Do Better

3.1 The Situative Perspective

Vygotsky's socio-cultural perspective Ideas central to the modern situative perspective were already prominent in the early 20th century writing of L. S. Vygotsky. Vygotsky claimed that mental abilities are grounded in available cultural tools and in social interaction (see [Wertsch and Tulviste, 1998] for a comprehensive discussion). When describing a psychological function, its genesis and development need to be understood. In terms of learning, what language, techniques, and examples were (and are) being used in acquiring a skill? What people were involved in the learning process, how did they interact, what were their backgrounds? Simply put, the nature of an individual's acquired knowledge is determined by the socio-cultural processes that have been shaping it (see [Wertsch and Tulviste, 1998]).

Modern situated theories Modern situative perspectives stress the social and cultural context or situation of learning. It is vital for the acquisition of a skill that it is taught in a situation similar to the intended situation of application. Rather than analysing individual agents, situative views focus on "interactive systems" ([Greeno, 1997], page 7) and people's participation in and experiencing of social situations. Research methods involve ethnographic studies and discourse analysis – quite in contrast to cognitive mathematical modelling. It is due to such fundamental differences that conflicts between the cognitive and the situative perspectives arise, with researchers dedicatedly arguing for their respective position as the superior one (follow the dispute of [Anderson et al., 1996], [Greeno, 1997], and [Anderson et al., 1997], with the later comment [Cobb and Bowers, 1999]).

Application of situated learning theories to practice calls for the embedding of learning material in real world contexts. It is therefore argued in [Seely-Brown et al., 1989] that students should be encouraged to learn by conversation in groups and the sharing of narratives. Abstract mathematical knowledge can be embedded in contextual stories meaningful to the learners and relevant for the application of the skills they learn. Transfer and generalisation would be obtained by reflecting on the conversations.

Example 3.1 (Palaver Tree Online) *The cooperative online learning environment in the domain of history "Palaver Tree Online" (PTO), presented in [Ellis and Bruckman, 2001] and [Ellis and Bruckman, 2002], is consistent with situative thought. This system is not an isolated piece of software. One learning situation involves a group of students who interview a contemporary witness online about his or her personal experience of a certain historical period (see Figure 2). The result of the learning process is a multimedia piece of life history composed by the students. Historical events are thus embedded in stories from real life. Students learn by cooperative work and interaction in a social situation. User models do exist in PTO, but those contain mainly* personal data *and* interests, *rather than knowledge states. Analyses of discourse and interaction are the main techniques in researching the learning process and designing the software.*

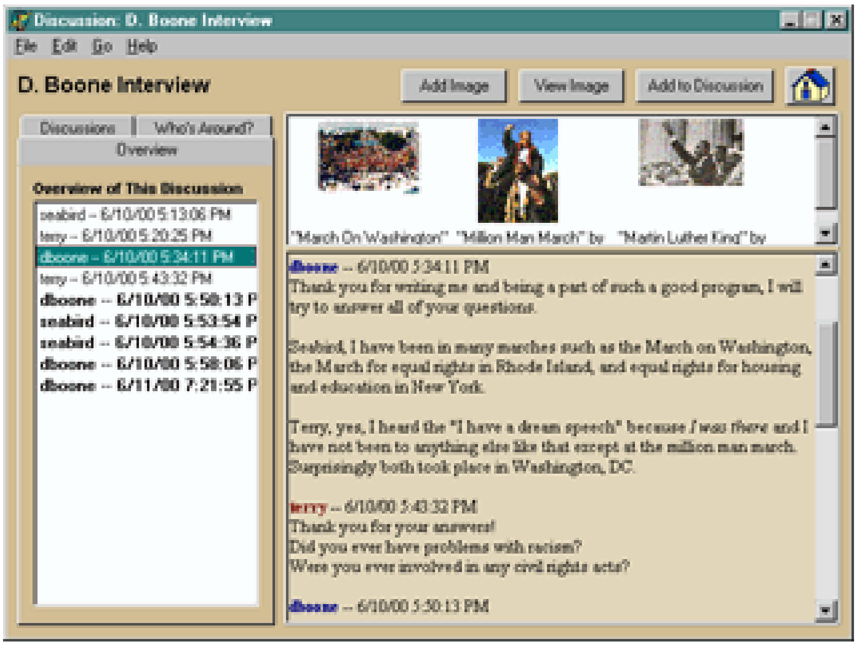

Figure 2: Palaver Tree Online

Indeed, the relationship between cognitivism and situativism appears to be dialectic in nature. The cognitive perspective of learning has its focus on the individual. The situative perspective focuses on group interaction. Cognitive teaching tends to happen "bottom-up", situated teaching "top-down". In Example 3.1, students would start with general information about the historical period, and then advance to a more detailed level, ie people's experiences of the time. Is history knowledge particularly suited for situative learning, and does maths knowledge suit the cognitive perspective? I would claim this is at most a tendency. "Bottom-up" and "top-down" learning are rather cognitive styles – *local/global* in terms of [Sternberg and Grigorenko, 1997], *serialistic/holistic* in terms of [Ford, 2000] – of individual learners in relation to a subject. Likewise, individual learning and group learning are student preferences – in terms of [Sternberg and Grigorenko, 1997] the *internal* and the *external* style. History knowledge can well be represented by a domain model, just with a different structure of prerequisite relationships than maths.

I agree with Greeno of the situative school that neither perspective will yield the ultimate theory of learning. He aims at shaping a synthesis starting with interaction processes and working towards the details of information structures ([Greeno, 1997], page 15). On the basis of situativism, he seeks to include cognitive ideas. For the improvement of ITS, I would aspire planning research the other way round:

> "One possible route takes the theory of individual cognition as its basis and builds toward a broader theory by incrementally developing analyses of additional components of situations that are considered as contexts for cognitive processes."
>
> [Greeno, 1997], pages 14 and 15

ITS improvement How can cognitive theories of ITS profit from situative thinking without loosing the precision of their models? Having a domain model as the framework for student modelling is a first step towards adding context to individual learning processes. In ITS, students solve practical domain problems experiencing "learning-by-doing" (cf [Anderson et al., 1996], page 8). Precise problems however may not be what learners encounter in real life. Problems in ALEKS were limited to "school culture", ignoring the complexity of real world situations. In a domain of history (Example 3.1), we can model objective knowledge about dates, political structures, and documented sources. We can however not grasp the variance in subjective accounts of more or less arbitrarily selected witnesses. Greeno writes:

> "In contrast to the behaviorist and cognitive views, where a domain of skills needs to be sampled, the situative view requires sampling across a domain of situation types in which participation involves the kinds of knowing that are of interest."
>
> [Greeno, 1997], page 8

As an improvement of the generic ITS in Figure 1, I suggest to replace problems with *situation types*. A situation type would comprise all "situations" in which a student can successfully or unsuccessfully apply the knowledge formerly needed to solve the "problem". When diagnosing a student's knowledge or teaching him or her, multiple instances of the situation type could be presented. Multiple examples increase the likelihood of transfer (see [Anderson et al., 1996]). ALEKS does use multiple examples, though mostly in similar and abstract types of situation. The situated examples proposed here would involve "the kind of activities in which we want [the student] to be successful" ([Greeno, 1997], page 13). Even better, they would involve activities the student wants to succeed in her- or himself. When studying the civil rights movement, a history student might be interested in particularly women's experiences. A young student of maths could want to apply adding fractions to sports statistics. I would therefore supplement the student model with data on interests that could be passed as parameters into situation types, thus generating adapted instances of situations: Figure 3.

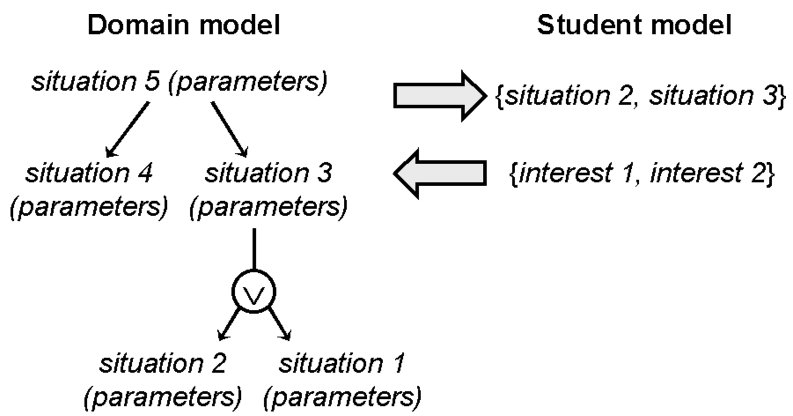

Figure 3: A generic ITS (2)

Vygotsky and the situative theorists stressed interaction. Increased interactivity – for a start with the system, maybe with other users – can be included in the exercises related to situations. There would be no need at this point to extend the meta-model for this cause.

3.2 Constructivism

Creative construction of knowledge There is a perspective of learning, originated from philosophy, which focuses on the individual and on practical context, but follows neither cognitive nor situative thinking. In *constructivism* (see [Säljö, 2000], Chapter 3 for its background), “a strong emphasis is placed on the learner as an active agent in the process of knowledge acquisition” ([de Jong and van Joolingen, 1998], page 179). Still, cognitivists distance themselves from the ideas (see [Anderson et al., 1995]). The crucial word in de Jong and van Joolingen’s definition is “active”. Applied to learning, constructivism means that one creatively constructs one’s knowledge oneself by experience. Instruction cannot result in learning without personal insight and discovery.

Scientific discovery learning The constructivist view of the world has inspired an innovative learning paradigm labeled *scientific discovery learning*. [de Jong and van Joolingen, 1998] is a review of the relevant literature. In such learning tools, a *simulation* of domain concepts is presented to the user. He or she is turned into a “researcher” who may manipulate the simulation and conduct experiments, in order to discover the underlying mechanisms. This view has certainly found its place in school, especially in the natural sciences – think of chemistry class.

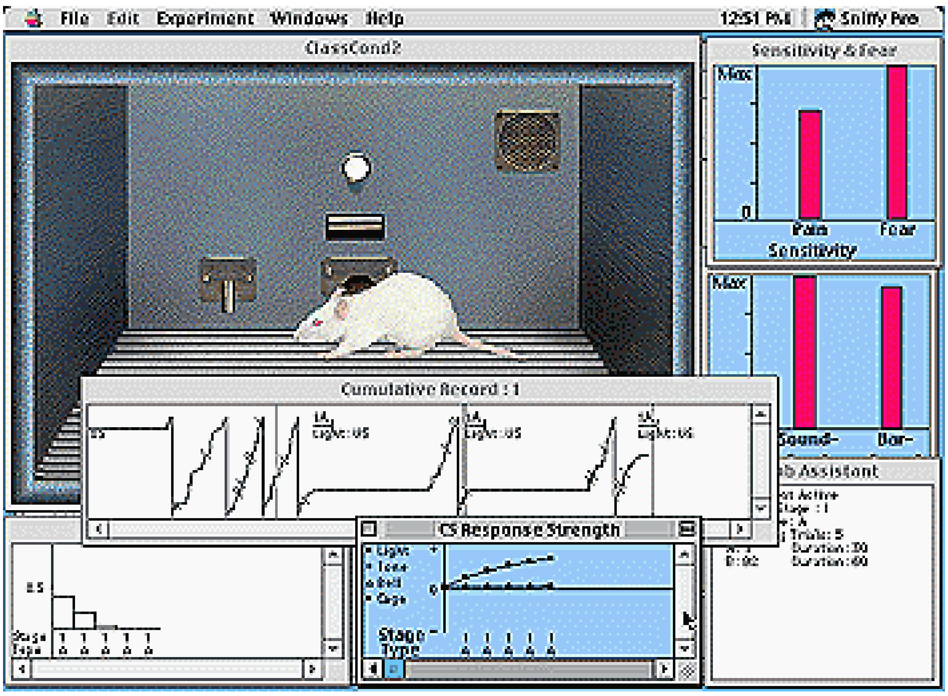

Figure 4: Sniffy the Virtual Rat

Example 3.2 (Sniffy the Virtual Rat) *This scientific discovery learning tool[2] provides a virtual laboratory for conducting experiments on classical and instrumental conditioning – as well as Sniffy himself, an animated laboratory rat who can be conditioned using lab inventory. The user finds a light, a speaker, a bar Sniffy can press, a water spout, and a food hopper for rewarding Sniffy with pellets. There are further instruments and diagrams to be checked, as well as a virtual lab assitant offering help in the form of text messages. See Figure 4. Without having tested the software, the concept sounds convincing. Simulated Sniffy looks great on screenshots of the virtual lab. Yet, "where is the human?" a situated theorist could ask, missing everyday life context.*

ITS improvement Constructivism seems compatible already with the intial meta-model of ITS (Figure 1). Exercises representing problems could be realised as experiments in a consistent lab setting. The other way round, Sniffy's lab assistant might be equipped with adaptivity (if it isn't already). The rat agent itself however should not adapt its behaviour to user knowledge, as that would negatively affect authenticity. Creative exercises need not be lab experiments. Simulation environments encouraging experimentation could include ALEKS' pencil, paper, ruler, and eraser, or in a programming domain an interpreter of the language in question.

3.3 Multidimensional student models

Adaptivity According to [Ragnemalm, 1999], Chapter 2, adaptation to student knowledge or skills is the usual objective in the design of ITS – which is sensible in the field of learning, where users gain knowledge. Further candidate differences between individuals are learning styles and *motivation*. Above, I suggested also to take user interests into account.

Learning styles Many cognitve styles have been identified and applied to learning. Sternberg and Grigorenko's *mental self-government approach* is descibed in [Sternberg and Grigorenko, 1997]. The authors present a classification of many different styles utilising metaphors borrowed from politics. I discuss cognitive styles in [Maciuszek, 2002] by applying them to the ALEKS learning environment. Main findings were that ALEKS falls short of supporting several expected inter-individual differences in learning styles.

Motivation There are more aspects of motivation than flow (discussed in [Norman, 1993]). [Ragnemalm, 1999], Chapter 2 names competition, chal-

[2] `http://www.wadsworth.com/psychology_d/special_features/sniffy/index.htm`

lenge, curiosity, confidence, and control as candidates for modelling inter-individual differences. A learner with a highly competative personality might be stimulated by games and rewards for good performance. An additional aspect of high or low motivation would be aspired or not aspired self-realisation.

ITS improvement I would call student modelling of only knowledge "one-dimensional". "Multidimensional" student modelling could further include significant interests, aspects of high motivation, and preferred learning styles: Figure 5. I already pointed out solutions for making ITS sensitive to new interests and situations of application, thus supporting also the *liberal* cognitive style. Inviting the user to participate in creative exercises would support *legislative* individuals. Preference for *visual* or *verbal* styles (cf [Oberlander et al., 1996]) could be adapted to by offering alternative ways of explaining wherever sensible in light of the content. Further styles and motivation could be handled by the design of exercises and feedback.

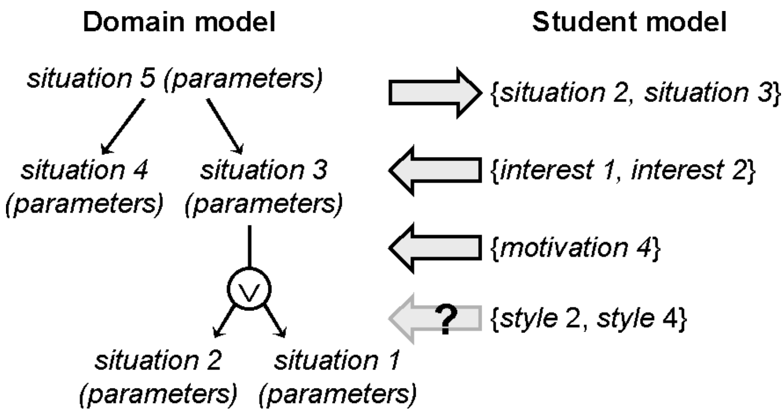

Figure 5: A generic ITS (3)

The practical evaluation found it a major drawback that ALEKS does not support "top-down" learning, referred to as the global style by Sternberg and Grigorenko. This trouble is not trivial. The "bottom-up" way of teaching may be an inherent feature of our modelling of domains. Greeno writes:

> "When a learning sequence is considered as a trajectory of skill/knowledge acquisition, it is natural to arrange activities that are systematic from the beginning, which will progress from being very simple to being more complex in line with analyses of behavioral prerequisites . . . and with processes of chunking in cognitive analyses. . . There are alternative ways to progress through the space. Early activities can be

> more complex and less systematic, so that as students progress their behavior becomes more systematic."
>
> [Greeno, 1997], page 11

Consider the graph in the meta-model. Arrows point only in one direction. Students cannot start with the top problem/situation and then work downwards. Allowing bidirectional arrows would require a rewriting of knowledge space theory. [Ford, 2000] reports of neural network approaches to the problem of adaptation to "bottom-up" (serialistic) and "top-down" (holistic) styles. Yet, this would introduce uncertainty into the exact knowledge structure of relations between items. I see four alternative and simpler solutions to the dilemma:

- Explanations to bottom-level exercises, especially such that introduce a new area, could give an additional high-level overview.
- The ALEKS pie chart representing the student model could give a more comprehensive overview of the domain.
- Although knowledge space theory does not allow symmetric relations (bidirectional arrows) between items, it does have alternative prerequisites (depicted as an "$\vee$" node in the figures). Style-based differences in paths through the graph might be modelled by alternatives.
- There could be two parallel knowledge structures. The first would be the conventional "bottom-up" structure. The second would arrange knowledge consistent with Greeno's "alternative ways". Teaching could use both for deriving two separate student models, or use one of them after a diagnosis of the student's preferred style.

4 Conclusions

Recapitulation The essay discussed conceptual strengths and weaknesses of intelligent tutoring systems in the light of modern theories of learning as well as kinds, levels, and styles of knowledge. ITS based on knowledge space theory were found to follow cognitive perspectives of learning and adapt to student knowledge. It was tried to form a synthesis of normally conflicting perspectives – cognitive and situative – in order to improve the educational value of ITS software. Suggestions were made to shift their focus from problems to situation types and from one-dimensional to multidimensional student models. Constructivist views were drawn on to argue for the design of learning environments that encourage experimentation and creativity.

Perspectives My current research interest is in intelligent software to help older people in everyday life situations. While such systems will involve computer based teaching in one form or another, the software paradigm will be more general than in this paper. Working interdisciplinarily between the social sciences and Computer Science, I have become very aware of conflicts involving the analytical "individual user" perspective versus the holistic "socio-cultural context" perspective. This essay provided some thoughts on how to bridge the gap approaching socio-cultural, situative context from an individual, cognitive perspective. I believe my ideas will generalise from the field of ITS to other knowledge based systems that employ "user models" instead of "student models" and "world models" instead of "domain models".

References

[Anderson et al., 1995] Anderson, J. R., Reder, L. M., and Simon, H. A. (1995). Applications and Misapplications of Cognitive Psychology to Mathematics Education. Unpublished manuscript, accessible on the Web, `http://act.psy.cmu.edu/personal/ja/misapplied.html`.

[Anderson et al., 1996] Anderson, J. R., Reder, L. M., and Simon, H. A. (1996). Situated Learning and Education. *Educational Researcher*, 25(4):5–11.

[Anderson et al., 1997] Anderson, J. R., Reder, L. M., and Simon, H. A. (1997). Situative Versus Cognitive Perspectives: Form Versus Substance. *Educational Researcher*, 26(1):18–21.

[Cobb and Bowers, 1999] Cobb, P. and Bowers, J. (1999). Cognitive and Situated Learning. Perspectives in Theory and Practice. *Educational Researcher*, 28(2):4–15.

[Craik and Lockhart, 1972] Craik, F. I. M. and Lockhart, R. S. (1972). Levels of processing. A Framework for Memory Research. *Journal of Verbal Learning and Verbal Behaviour*, 11:671–684.

[de Jong and van Joolingen, 1998] de Jong, T. and van Joolingen, W. R. (1998). Scientific Discovery Learning With Computer Simulations of Conceptual Domains. *Review of Educational Research*, 68(2):179–201.

[Doignon and Falmagne, 1999] Doignon, J.-P. and Falmagne, J.-C. (1999). *Knowledge Spaces*. Springer-Verlag, Berlin.

[Dreyfus and Dreyfus, 1986] Dreyfus, H. L. and Dreyfus, S. E. (1986). *Mind over Machine*, chapter 1, pages 16–51. Basil Blackwell Ltd, Oxford.

[Ellis and Bruckman, 2001] Ellis, J. B. and Bruckman, A. S. (2001). Designing Palaver Tree Online: Supporting Social Roles in a Community of Oral History. In *Proceedings of ACM CHI 2001 Conference on Human Factors in Computer Systems*, Seattle.

[Ellis and Bruckman, 2002] Ellis, J. B. and Bruckman, A. S. (2002). What Do Kids Learn from Adults Online? Examining Student-Elder Discourse in Palaver Tree. In *Electronic Proceedings of CSCL 2002 Conference on Computer Supported Cooperative Learning*, Boulder.

[Ford, 2000] Ford, N. (2000). Cognitive Styles and Virtual Environments. *Journal of the American Society for Information Science*, 51(6):543–557.

[Greeno, 1997] Greeno, J. G. (1997). On Claims That Answer the Wrong Questions. *Educational Researcher*, 26(1):5–17.

[Maciuszek, 2002] Maciuszek, D. (2002). Cognitive Evaluation of the Maths Tutoring System ALEKS. Course paper "Learning, Cognition, and IT" (Part 1).

[Norman, 1993] Norman, D. (1993). *Things that Make Us Smart*, chapter 2. Addison, Wesley, Reading.

[Oberlander et al., 1996] Oberlander, J., Cox, R., and Monaghan, P. (1996). Individual differences in proof structures following multimodal logic teaching. In *Proceedings of the 18th Annual Conference of the Cognitive Science Society (COGSCI)*, pages 201–206.

[Ohlsson, 1995] Ohlsson, S. (1995). Learning to Do and Learning to Understand: A Lesson and a Challenge for Cognitive Modelling. In Reimann, P. and Spada, H., editors, *Learning in Humans and Machines: Towards an Interdisciplinary Learning Science*, pages 37–62. Pergamon Press, New York.

[Ragnemalm, 1999] Ragnemalm, E. L. (1999). *Student Modelling based on Collaborative Dialogue with a Learning Companion*. Dissertation, Linköping University, IDA.

[Seely-Brown et al., 1989] Seely-Brown, J., Collins, A., and Duguid, P. (1989). Situated Cognition and the Culture of Learning. *Educational Researcher*, 18(1):32–42.

[Säljö, 2000] Säljö, R. (2000). *Lärande i praktiken: Ett sociokulturellt perspektiv*. Prisma, Stockholm. In Swedish.

[Sternberg and Grigorenko, 1997] Sternberg, R. J. and Grigorenko, E. L. (1997). Are Cognitive Styles Still in Style? *American Psychologist*, 52(7):700–712.

[Tulving, 1985] Tulving, E. (1985). How Many Memory Systems Are There? *American Psychologist*, 40(4):385–398.

[Wertsch and Tulviste, 1998] Wertsch, J. V. and Tulviste, P. (1998). L. S. Vygotsky and contemporary developmental psychology. In Faulkner, D., Littleton, K., and Woodhead, M., editors, *Learning Relationships in the Classroom*, pages 13–30. Routledge, London.