Design and Evaluation of an Apprenticeship Setting for Learning Object-Oriented Modeling

Jakob Tholander, Klas Karlgren, Fredrik Rutz, Robert Ramberg

Department of Computer and Systems Sciences, Stockholm University,
E-mail: {jakobth, klas, rutz, robban}@dsv.su.se

This paper presents a study of three different tools for learning object-oriented analysis. The tools are built on the cognitive apprenticeship learning model, thus emphasizing practical, authentic skills and realistic problem solving. Object-oriented analysis is a typical design task in which solutions to a problem may differ greatly but still be equally valid and we therefore find the cognitive apprenticeship learning model to be particularly suitable. 20 students participated in the study and performed a modeling task with one of three learning tools. A conclusion is that the learning tools support learning of the skills that are central to expert modelers. We also found that learners are used to a traditional supervised learning style and therefore need support to use the tools in order to develop independent problem solving skills.

Keywords: **Cognitive apprenticeship, pedagogical foundations, system eval.**

1. Introduction

Situated learning theories and especially the cognitive apprenticeship model have reacted against the overemphasis on abstract knowledge of many instructional models and have instead shifted the focus of learning and instruction towards practical, authentic skills and realistic problem solving.
We use cognitive apprenticeship as a starting point for designing computer based training material for object oriented (OO) modeling [12], [6]. OO models are constructed during information systems development to build a common view and language for decision making and to serve as a basis for the design and implementation of a computer system [1]. Often the teaching of OO modeling has an emphasis on the concepts and theory behind OO models. Usually courses and textbooks do not focus much on the practical skills concerned with the actual modeling process, i.e., the constructing of the model. Cognitive apprenticeship is argued to help learners to focus on the activities that are central to the behavior of expert problem solving and to the development of such skills [4]. Those activities can be divided into firstly meta-cognitive activities like planning, problem identification, goal definition, and evaluation, and secondly into activities that give access to how experts and other peer learners solve problems. In this paper we report a study of three tools for the support of apprenticeship style learning and how each tool affected students' reasoning and problem solving.

2. Cognitive Apprenticeship

Collins *et al.*, [4], propose a framework for developing learning environments based on cognitive apprenticeship. In this work we have focused on the following aspects of that framework:

1. *Modeling and observation*. One important aspect of apprenticeship learning as described by Lave & Wenger, [7], is the possibility for learners to get access to how experts solve real world problems. By observing experts, learners can pick up tacit knowledge in order to develop awareness of what they need to learn and how different skills are used.
2. *Reflection, evaluation, and planning*. Collins *et al.*, [4], emphasize the importance of meta-cognitive skills in the development of expertise. Learners need to be encouraged to plan their problem solving, reflect over and evaluate their own and other solutions as well as the status of their own knowledge.
3. *Practice using authentic tools*. The situated learning perspective emphasizes that knowledge and skills cannot be separated from the context and community where it is used and that knowledge always involves practical activity. Therefore learners need to be trained by practical activities in an environment in which all the features of the real context are present and preferably with experts and other peer learners.
4. *Articulation and use of expert language.* The situated perspective emphasizes that knowledge is tightly coupled to participation in communities of practice and to the social interaction in those communities. Therefore learners need to be exposed to the way practitioners of those communities carry out activities and be given possibilities to practice using the authentic

language – the language game - and to articulate their knowledge in those communities [10].

3. The Domain: Learning Object Oriented Analysis

The purpose of OO analysis is to create models of an enterprise that can be used for the subsequent design and construction of supporting communication and information systems. The models should preferably also be designed as general and reusable components realizing the idea of patterns. When students learn to construct such models, it is essential they get the opportunity to work with problems of a realistic size and complexity. Working with large problems may, however, be overwhelming for inexperienced students. Therefore, extensive and qualified supervision is needed to assist students in their work. Furthermore OO analysis is a typical design task in which solutions to a problem may differ greatly but still be equally valid. The traditional in-school modeling situation differs extensively from the situations in which professionals act. Professionals are often confronted with unknown domains which are hard to understand. This is seldom the case in textbooks and school exercises. The skills involved in such modeling activities are hard to make explicit and the cognitive apprenticeship model is a promising way of supporting the learning of such skills [6], [12].

4. Design of Learning Tools for Modeling Apprenticeship

In order to provide learner support (or scaffolding) for the aspects of cognitive apprenticeship that we wanted to focus on we have designed and implemented the tools illustrated in Figure 1 through 3 below. These were used together with a commercial design tool for creating OO models.

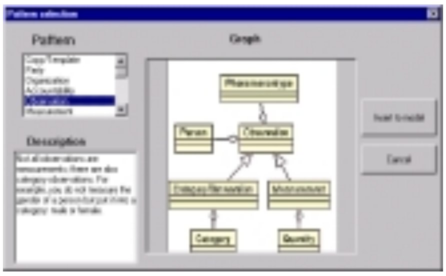
Fig. 1: The Pattern library


Fig. 2: Expert track


Fig. 3: The assistant

A Pattern Library

In the object-oriented community within system development the notion of patterns has become important. A modeling pattern is a prototypical or abstracted OO model, which can be applied in different real-world situations with appropriate modifications [5]. For instance, patterns like product, ordering, booking, and contract may be reused in many order management systems. One might oppose to focusing on abstractions when adopting a situated view on knowledge. But, abstract knowledge is not less situated, since abstraction in itself is a social activity, with specific goals and customs, and with certain expectations about success of the abstraction. Our goal here is thus to support students to make abstractions, not to learn the specific abstractions [2]. Therefore the pattern library is closely integrated with the modeling tool and it is possible to insert patterns into the model currently worked with.

We also see the pattern library as a way of supporting learners to practice use of, and articulation in, the particular language and concepts used by the practitioners of a field - the language game [10]. Similarly, Morrison & Collins [9], discuss learning in terms of becoming fluent in epistemic games which are often language-based. The pattern library supports aspects 2, 3, and 4 of the framework discussed in section 2 for cognitive apprenticeship learning environments.

A Critical Assistant to Support Reflective Thinking

At first sight, it might seem improbable that an automated tool should be able to criticize OO models, as this is a most qualified task. Assistants have been used in learning environments to provide a look and feel of a tutor [8]. It is our experience from supervising students that deep and precise critique of their models is not necessary. In fact, it is usually preferable to give only a more shallow critique, asking questions such as "Why did you introduce this object?" or "Are you able to represent this state of affairs in the enterprise?" Asking such questions forces students to reflect on their models and explain them to themselves, which is a most effective instrument in a learning situation [3]. Constructing this type of shallow questions and comments is a task that is well suited for a computer-based tool. We use three categories of questions: i) specific questions about some of the student's constructs, for instance, “Are you aware of the fact that you have a many-to-many relation between the object ‘Person’ and object ‘Flight’?” ii) general questions about how the student intend to handle some part of the domain description, for instance, “Make sure that your model can handle that airplanes might be delayed or rescheduled to another destination?”. iii) general comments to encourage reflection and critical thinking. For instance, “Are you sure that is a good idea?” or “Was that well considered”. The critical assistant support aspect 2 of section 2.

Expert Problem Solving Tracks

The third tool consists of expert tracks, i.e., step-by-step accounts of how experienced modelers proceed through modeling problems showing how the modelers reason, solve problems, and make decisions during the modeling process. To reach a high degree of authenticity in these tracks we have video-recorded experienced modelers modeling similar problems. The purpose of recording the modelers was to provide examples of possible ways of how experts solve modeling problems rather than exactly pinpointing the general strategy of modeling in a definite way. From the recordings, the modeling sequence and especially critical points in the problem solving process were extracted and presented in the tracks (in the form of hyper-linked web pages). At each step questions and unclarities are made explicit, and users have to click their way on to find out the expert's decision and motivation about how to proceed. Focusing on the decision-making of experienced modelers hopefully encourages reflection and articulation on the part of the user. From these authentic steps performed by an expert solving a problem students can observe expert behavior with all its particulars [2]. Since the tracks are designed to include the modeling of different modelers they will provide multiple perspectives on the same problem, as discussed by [4], [11]. The expert tracks supports aspects 1, 2, and 4 of section 2.

4.1. Hypotheses and Goals of the Study

The general hypothesis of this study was that subjects using our tools produce OO models in a manner closer to how experts work. More specifically we believed that the different components would help the students to spend more time on meta-cognitive and high-level activities, and also that they would create component-like models. One goal of the study was to investigate if and how each component affect the way students solve modeling problems and if and how this affect their final results. We also wanted to gather usability data on our designs and implementations. For each component we had the following hypothesis for how the students would use or react to it: Firstly we believed that the students would use the pattern library in difficult situations or when confused and did not know how to continue. Secondly we expected the assistant to help subjects spend more time and effort to critically reflect upon their model and evaluate their solutions. Finally we expected subjects using the expert track to refer to the expert's reasoning when solving their own problems.

5. Method and Design of the Study

Since the view on learning adopted here emphasises the way activities are carried out rather than the outcome of the activities, our goal has been to achieve a design where we focus on the learning process rather than the results. Twenty subjects participated and worked in pairs using the tools described above. All subjects were computer and systems science students with some experience of OO modeling. The students were given a modeling task concerning an example about terrorists and airplane hijackings. The modeling domain was described in a text and the subjects’ task was to create an OO model in a modeling tool. In order to get information about how and if the students’ reasoning were affected by the software we let the students work in pairs of two on a, for them, fairly difficult modeling problem. OO modeling is typically a social activity, done in groups or in seminars, so we let the subjects work in pairs, which also encouraged them to discuss their reasoning and giving a think-aloud effect. There were ten pairs, two were given the pattern library, two the expert track, two the assistant, two used all tools, while two pairs modeled without help from any tool. As the number of subjects was small, we did not make any quantitative analyses.

The maximum time for solving the problem was 45 minutes except for the group with all the components, which could use 75 minutes. We deliberately gave them a problem that would be too hard for them to complete in the time they were given so that it would be challenging for all groups. Each group was recorded on video from the front so we could see their faces and interactions and from the back so we could see the computer screen and where on the screen they pointed.

Before the modeling task we asked the students about their opinions about modeling tools in general, and in particular about tools for learning how to model, and about the different approaches we have described in this paper. After the modeling the students filled in a questionnaire and we concluded with an interview asking the same questions as we had asked before the modeling and we discussed how they had experienced the modeling session and the tool they were given. As we handed out the questionnaire we asked the students to save any comments for the concluding interview, as we had not left any space for comments on the questionnaire, instead promoting discussion during the final interview.

6. Results

Below, picture 4 and 5, are two examples of the models created by the groups. The models were constructed by group 4 (Figure 4) and by group 6 (Figure 5) and show the quality-span and difference in character between the different solutions.

To examine whether the different tools were used and whether they affected the subjects' modeling, three experienced modelers were asked to predict which tool the subjects had used during their modeling. We found that it is fairly easy for an expert to spot which groups had access to the pattern library. Probably since many of the concepts from the patterns appeared in the models. Most predictions concerning the tracks were correct which indicates that the students referred to the tracks when solving their problems. Note that these predictions do not have anything to do with the quality of the models, only that the tools did effect their work. The quality aspect is discussed below.

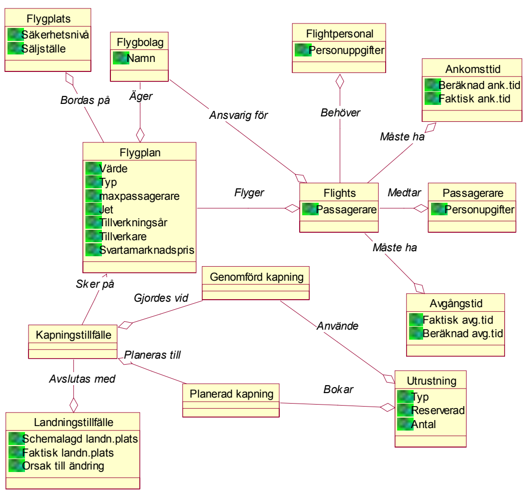

Fig. 4: An example of a component-like and rather complete model

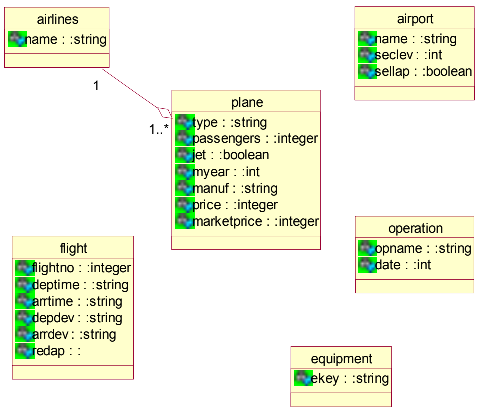

Fig. 5: An example of detailed but not so complete model

6.1. Different Modeling Strategies: Jumpy or Chunky

The subjects had different modeling styles. Some subjects had a rather linear modeling style; they read through the whole text and extracted classes mainly by choosing the nouns. One of the subjects in group 4 started modeling by exclaiming "maybe we should look for objects then!". Another pair (9) first drew all the classes they needed, and later started adding relations between them

The subjects modeling styles could also be described as more or less *jumpy* and more or less *chunky*. The modeling of subjects who moved very much between different parts of the OO model was described as jumpy. Others worked in a less jumpy manner by starting with one part, completed it, and then moved on to the next most closely related part.

A chunky modeling style means that the modeling focuses around logical parts of the model that are larger than one class at the time; often whole modeling patterns. A chunky modeling style may, e.g., involve reasoning about resources, resource allocation, resource planning, and resource use as a whole "chunk" of classes. Perhaps the more experienced modelers would start modeling at chunkier level and move towards more detail. Experienced modelers will typically use a chunky and jumpy modeling style. They may use different modeling patterns and usually start by laying out the bigger chunks, and move back and forth between the different parts of the entire model getting into more detail. They would start in the lower right corner of Figure 6 and move up and finish in the top right corner. From Figure 6 one can see that groups 3 and 5 had the most chunky modeling process and especially group 3 jumped between chunks while modeling. Groups 2 (assistant), 4 (all tools) and 9 (no tools) had a tendency to jumpy modeling processes but not very chunky. Group 6 which used the assistant were neither chunky nor jumpy in their modeling process, i.e., they proceeded in a mostly linear fashion. Values from 0 to 6 on the two dimensions in Figure 6 were given for each group by analyzing the video recordings (group 1 and 8 has not been analyzed yet).

One can see that we have not characterized many of the groups as having a chunky modeling process, most groups jump around between smaller parts in their modeling process. Since most of the groups did not finish their assignments, we would have preferred to have had more groups further out on the chunky dimension, and not as high in the jumpy-dimension.

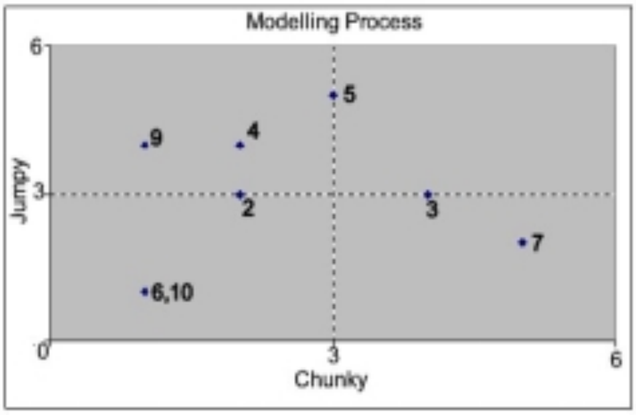

Fig. 6: Characterization of the dimensions of chunkiness and detailedness.

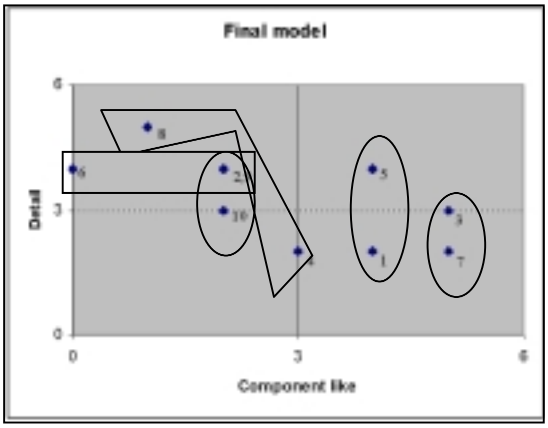

Fig. 7: Level of detail and use of components in the final models. Groups using the same tool are circled.

6.2. Component-like and Detailed Models

The resulting models can be described as more or less detailed and more or less component-like. A model which is component-like consists of logical parts which resemble the idea of patterns. A model which is detailed consist of many attributes, relations and classes. These are not two exclusive dimensions. A complete model should both be component-like and detailed. In the case of less complete models a component-like model is preferred before a detailed model since it is easier to extend and change a component-like model than a model which is only detailed. Values from 0 to 6 on the two dimensions in Figure 7 were given for each group by analyzing the final models.

From Figure 7 above one can see that the final models of the groups that used only the pattern library and the groups that used only the tracks were more component-like, or pattern based. One can also see that the groups that only used the assistant created models with a higher level of detail.

The final models of the pairs who used all tools were not of an especially component-based character and this could be viewed as a failure. If the patterns and the track was support for the pairs using only these, why were they not a help for the pairs of subjects who could use all tools? One interpretation could be that using all tools is more time-consuming which would explain why these pairs in the limited time available could not produce as good a model as the others. Another possible interpretation is that providing all tools is confusing since the tools were not integrated with one another, and that the subjects therefore do not pay as much attention to these as compared to if only one tool is provided.

6.3. Evaluations of the Final Models

When the subjects had finished modeling they were asked to evaluate how good they thought their model was on a scale 1-6 where "1" is "very bad" and "6" is "very good". The subjects' own ratings of their final models corresponded rather well with the expert ratings. The subjects who modeled with the assistant, however, seemed to slightly overrate their models as compared to the expert evaluations. We also noticed that the experts graded the groups which had used the pattern library only and the tracks only higher than the rest of the groups. Further there seem not to be much difference between how the subjects in each group rated their own models.

6.4. Expectation Ratings and Ratings of the Quality of Information

After the subjects had finished modeling, the subjects received a questionnaire with questions concerning the modeling and the tool(s) they used. Among many other questions, they were asked to rate how well the tool(s) they used fulfilled the expectations they had on the tool, the quality of the information in the tool, whether the tool supported reflection about the model, and whether the tool helped during their modeling.

The subjects' ratings of the quality of information in the tools corresponded well with ratings of how well the tools fulfilled their expectations of the tools. The subjects' expectations on the tools were better fulfilled for the pattern library and the track as compared to the assistant. They also rated the quality of information higher for the pattern library and the track than the assistant.

7. Discussion

7.1. Was the Expert Track of Any Help?

The hypothesis for the expert track was that the subjects would refer and relate to the expert's reasoning when solving their own problems. This can be done in different ways, either by explicitly discussing the expert's solutions when modeling, or more implicitly using the same kind of solutions. For instance, pair 4 returned to the track and compared their modeling to the modeling of the experts model and tried to see analogies between the solutions. They tried to directly map the classes of a solution in the expert's model to the concepts in the domain they were trying to model.

One of the pairs (no. 3) explicitly said they did not want to use the same kind of solution as the expert had because they thought they had found a better one themselves. In a way this is a kind of relating to the expert's reasoning although they disagreed with the expert, and hopefully this lead to more reflection on their modeling. These two subjects were rather experienced modelers, and they were also the most critical towards the expert track for being too easy. This pair discussed the resource allocation-pattern used in the track and pointed to its limitations and considered their own solution more advanced. A reasonable conclusion from this is that expert tracks have most value for learners who are not yet so proficient.

Moreover, the subjects who did not go through the track seemed often to run into problems which may have been alleviated with the expert track. One group that wasn't given the expert track had trouble getting started with the modeling. They ended up drawing a few objects but did not get to connecting these with any but one single relation. Giving the students the possibility to observe how an experienced modeler approaches an analogous problem would have been a help for this group to get started. Another pair (no. 9) which modeled without the help of any of the tools initially ran into some very fundamental questions concerning the task. They were unsure how much and what kind of information should be included in the model: "Do we need to model just what is here [points at the text]? Is the purpose that this should eventually be implemented ... because in that case perhaps GUI-classes should be included…". If these subjects could have been able to observe an experienced modeler solve analogous problems such questions would have been answered implicitly and they would not have had to have been asked by these subjects.

In the interviews several subjects said they would appreciate being able to observe an experienced modeler constructing models of analogous modeling problems. Also many of the subjects who followed the expert track said they were pleased with it. When asked what kind of support for modeling the subjects would like, one subject wished for more supervision and guidance during modeling. Another liked the idea of tracks because it may help "finding the concepts". Another subject wished for more examples. We believe it may be beneficial to provide expert tracks showing how experienced modelers solve analogous models and any limitations in their solutions.

7.2. Did the Students Consult the Pattern Library?

From the two groups that used only the pattern library one could see that the patterns had influenced the character of their final models. Many parts of their models contained concepts and constructs that resembled several of the patterns, or parts of the patterns in library. The patterns also helped them to keep their models at a higher level of abstraction. From the videotapes we could see that they often consulted the pattern library when they did not know how to continue. For instance, one subject said that "the patterns were good to use as a way to get new ideas and as a help to see things from a new perspective". One subject also commented that they looked at the patterns a lot but that they did not use them exactly as described in the library, but that they "stole" the ideas and adapted them to their own solution. This subject thought that they did not use the library as intended, but we believe that this is exactly the way the pattern library should be used. The students must create their own understanding and interpretation of patterns and component-based thinking. Patterns prescribe a way of looking at OO models as reusable and recurring components; patterns are not specific information that must be memorized.

One subject who worked with all three tools said, "I wish we had used the pattern library more since we got many good ideas from it, but we are not used to work that way". We interpret this as students are not used to search for what they need but are used to being "served with" the information they should learn. Another group who worked with all the tools said they liked the pattern library because "you could use it on your own initiative when you needed new information."

7.3. Did the Assistant Support Reflection and Critical Thinking?

From the final solutions one could not see any direct influence by the assistant but the subjects were definitely affected by the assistant's comments and statements during the modeling. Many times when the assistant commented about some particular class they just nodded or said "OK", but sometimes they said things like "we are not finished with that yet" or "yes, but we are working over right now". We believe that this indicates that the timing of the assistants statements are extremely important because mistimed statements risk being disturbing and cause the students to stop paying attention to what the assistant says.

When the assistant commented about some important aspect of the domain that needed to be covered in the model it sometimes had the effect we wanted to achieve, namely that the students would evaluate and reflect over their model. After the assistant had commented about the importance of considering how to model information about redirected flights one subject said: "What?, that depends on what we need to cover, we could do it with a method that changes the old time, but then we lose previous information about times etc, hmm". Such domain specific comments were sometimes a source of confusion particularly when the students were deeply involved in some aspect and the assistant commented on something not immediately relevant.

From different groups we could see that they wanted the ability to ask the assistant for help. One subject said while he and his partner discussed the possible use of date class "Maybe we could ask him" and another subject when she was confused said "Hey, give us some help here", and pointed to the assistant. The subjects' opinions about the assistant varied much between the groups. One group believed it provided useful information while another group thought it was in their way and that they wanted to search information on their own initiative instead of having it forced upon them by the assistant. One important factor in the opinions towards the assistant is, as stated above, how the subjects experienced the timing of the assistant's comments, but we also believe that this varies depending on the different learning styles of the subjects. Some subjects said "I want to work alone for a while and then ask a tutor" while other subjects seem to be more flexible for interruptions.

One observation concerning the groups that used all the tools was that the assistant had a tendency to take much attention away from the other two tools. One reason for this is that the assistant metaphor is a new and somewhat exciting form of interface, and also that it used audio to interact with the students. It also suggests that the assistant is needed to be closely integrated with the rest of the tools in order for them all to work efficiently.

7.4. General Observations about the Subjects' Modeling

One observation about many of the subjects modeling (independently of the tool they used) was that they often lost themselves on an overly detailed level of modeling. They often spent a considerable amount of time on discussing the attributes needed for particular classes, when it would have been much more fruitful to discuss the relations they needed in order to represent important aspects of the domain. Several groups also lost their focus and got engaged in aspects that was not central for solving the problem. For instance, one groups model consisted of only five classes but each with quite a few attributes, while their complete model only contained one relation between different classes. Such students, would have benefited with strategic scaffolding that is initiated by the system (the assistant), instead of scaffolding that is accessed only when the students realize their need for strategic support as is the case with the expert tracks and the pattern library.

We also gave the same problem to a one of the teachers at our department. She finished the problem in about one hour and her solution was rather complex. While several of the students considered their solutions to be good or fairly good their models did not cover as much information as the teacher's model and the students had missed many of the problematic issues in their solutions. The problem was not that the students struggled with trying to solve those issues without succeeding, but that they did not see that they were problematic. This indicates that an important aspect of expertise in this domain is to be able to see the total span of a problem and to be able to identify important issues in the problem domain. From this we conclude that students need to be provided even stronger support that helps them develop skills for problem identification and evaluation of solutions than they had with the tools they were given.

8. Conclusions and Future Work

The pattern library and the expert tracks are typical constructivist learning tools which the students can use as they want, thus supporting them in creating their own understanding of things and help them search for the information needed to solve a particular problem. Such tools are good for students that already have developed some of these skills which were the case for several of our subjects. But some of the subjects did not use the library nor the tracks as much and would have benefited with help and suggestions from the system on when to use these tools. This was indicated by the interviews. Further found that the students did not go back to the expert tracks to get help as much as we wanted, but the tracks still influenced the students' solutions. The assistant would be useful for providing pointers to patterns, expert solutions, and also to other information that could be useful in particular situations. Some of the subjects who used all three tools also commented that the assistant would be beneficial for this. A common comment about the assistant was that assistants is very useful as long as it does not say the wrong things and that it does not say it at the wrong moment. Further, even though the pattern and the expert track groups did better than the assistant groups we do not want to conclude that it is a worse tool to learn to model with, but rather that the pattern library and the expert tracks provide more practical help for problem solving. We believe the assistant provides support of a more strategic nature which gives long term benefits and must be combined with other tools. The shallow support given by the assistant was not always successful. This must not falsify our hypothesis but may indicate failures of mistimed interruptions.

Central to our future research activities is to integrate the different components, since we observed that the groups that worked with all three components got confused and sometimes had problems to focus on the modeling task itself. The assistant will be the central part from which the other components can be used or pointed at. We find the assistant metaphor to be a very promising way to provide scaffolding that is initiated by the system. We therefore work on having the assistant making suggestions, on its own initiative, about where in the pattern library or in the expert tracks the student could find useful support for some particular problem. Furthermore we will also make it possible for the students to initiate interaction with the assistant by letting students ask it for support. The assistant could then provide hints and suggestion about where in the pattern library the expert tracks to look, or by just commenting their model. As students get more advanced and gain insights about the knowledge or skills they lack or possess, we could fade the support given by the assistant by decreasing the amount of interruptions.

References

[1] Boman, M., Bubenko, J., Johannesson, P., & Wangler, B. (1996). *Conceptual Modeling*. (6th ed.). Stockholm: Department of Computer and Systems Sciences, KTH and Stockholm University.

[2] Brown, J. S., Collins., A., & Duguid, P. (1989). Situated Cognition and the Culture of Learning. *Educational Researcher, 18*(1), 32-42.

[3] Chi, M. T. H., de Leeuw, N., Chiu, M.-H., & Lavancher, C. (1994). Eliciting Self Explanations Improves Understanding. *Cognitive Science, 18*, 439-477.

[4] Collins, A., Brown, J. S., & Newman, S. (Eds.). (1989). *Cognitive Apprenticeship: Teaching the Crafts of Reading, Writing, and Mathematics*. New Jersey: Lawrence Erlbaum Associates.

[5] Fowler, M. (1997). *Analysis Patterns: Reusable Object Models*: Addison Wesley Publishing.

[6] Karlgren, K., Tholander, J., Dahlqvist, P., & Ramberg, R. (1998). *Authenticity in Training Systems for Conceptual Modelers.* International Conference of the Learning Sciences, Atlanta, Georgia.

[7] Lave, J., & Wenger, E. (1991). *Situated Learning: Legitimate Peripheral Participation*: Cambridge University Press.

[8] Lester, J. C., et al. The Persona Effect: Affective Impact of Animated Pedagogical Agents. At Chi'97 (1997).

[9] Morrison, D., & Collins, A. (1995). Epistemic Fluency and Constructivist Learning Environments. *Educational Technology,* 39-45.

[10] Ramberg, R., & Karlgren, K. (1998). Fostering Superficial Learning. *Journal of Computer Assisted Learning, 14*, 120-129.

[11] Spiro, R. J., Feltovich, P. J., Jacobson, M. J., & Coulson, R. L. (1992). Cognitive Flexibility, Constructivism, and Hypertext: Random Access Instruction for Advanced Knowledge Acquisition in Ill-structured Domains. In T. M. Duffy & D. H. Jonassen (Eds.), *Constructivism and the Technology of Instruction* . Hillsdale, NJ: Lawrence Erlbaum.

[12] Tholander, J., Karlgren, K., Dahlqvist, P., & Ramberg, R. (1998). *Cognitive Apprenticeship in Training for Conceptual Modeling.* Paper presented at the WebNet98, Orlando, Florida, USA.