Adaptive Navigation Support in Educational Hypermedia: the Role of Student Knowledge Level and the Case for Meta-Adaptation

Peter Brusilovsky

School of Information Sciences

University of Pittsburgh

Pittsburgh PA 15260

Email: peterb@mail.sis.pitt.edu

Adaptive hypermedia is an alternative to the traditional "one-size-fits-all" approach in the development of hypermedia systems. Adaptive hypermedia (AH) systems build a model of the goals, preferences and knowledge of each individual user, and use this model throughout the interaction with the user, in order to adapt to the needs of that user. Adaptive navigation support is a specific group of adaptive hypermedia techniques that become especially popular in educational hypermedia systems. This paper provides a brief overview of main adaptive navigation support techniques and analyzes the results of most representative empirical studies of these techniques. It demonstrates an evidence that different known techniques work most efficiently in different context. In particular, the studies summarized in the paper have provided evidence that users with different knowledge level of the subject may appreciate different adaptive navigation support technologies. The paper argues that more empirical studies are required to help the developers of adaptive hypermedia systems in selecting most relevant adaptation technologies. It also attempts to build a case for meta-adaptive hypermedia systems, ie, systems that are able to adapt the very adaptation technology to the given user and context.

1 Introduction

Adaptive hypermedia is an alternative to the traditional "one-size-fits-all" approach in the development of hypermedia systems. Adaptive hypermedia (AH) systems build a model of the goals, preferences and knowledge of each individual user, and use this model throughout the interaction with the user, in order to adapt to the needs of that user (Brusilovsky, 1996). For example, a student in an adaptive educational hypermedia system will be given a presentation that is adapted specifically to his or her knowledge of the subject (De Bra, & Calvi, 1998), and a suggested set of most relevant links to proceed further (Brusilovsky, Eklund, & Schwarz, 1998).

As was pointed out in (Brusilovsky, 1996), AH systems can be useful in any application area where a hypermedia system is expected to be used by people with different goals and knowledge and where the hyperspace is reasonably big. Users with different goals and knowledge may be interested in different pieces of information presented on a hypermedia page and may use different links for navigation. AH tries to overcome this problem by using knowledge represented in the user model to adapt the

information and links being presented to the given user. Adaptation can also assist the user in a navigational sense, which is particularly relevant for a large hyperspace. Knowing user goals and knowledge, AH systems can support users in their navigation by limiting browsing space, suggesting most relevant links to follow, or providing adaptive comments to visible links. It is quite naturally that educational hypermedia was one of the first application areas for AH. In educational context, users with learning goals and knowledge on the subjects require essentially different treatment. It is also in educational hypermedia where the problem of "being lost in hyperspace" is especially critical. A number of pioneer adaptive educational hypermedia systems were developed between 1990 and 1996.

2 Empirical Studies of Adaptive Navigation Support in Educational Hypermedia

2.1 Adaptive Navigation Support

There are a number of different methods and techniques that are used in adaptive educational hypermedia (Brusilovsky, 1996). Among others, a group of techniques known as adaptive navigation support become especially popular in adaptive educational hypermedia. The idea of adaptive navigation support techniques is to help users to find their paths in hyperspace by adapting link presentation to the goals, knowledge, and other characteristics of an individual user. Adaptive navigation support can guide the students both directly and indirectly and can work with much larger amounts of learning material using much simpler student models. In a WWW context where hypermedia is a basic organizational paradigm, adaptive navigation support can be used naturally and efficiently. There are several known ways to adapt the links. The most popular technologies are *direct guidance*, *sorting*, *hiding*, *annotation*, and *generation* (Brusilovsky, 2001).

In the Web-based education context, direct guidance and adaptive annotation was pioneered in ELM-ART (Brusilovsky, Schwarz & Weber, 1996; Weber & Brusilovsky, 2001) and since then applied in descendants of ELM-ART such as InterBook (Brusilovsky, Eklund & Schwarz, 1998) and numerous other systems. Figure 1 shows direct guidance and adaptive annotation in InterBook (Brusilovsky & Pesin, 1998). Adaptive hiding and disabling was pioneered in 2L670 (De Bra, 1996). Adaptive link generation is the newest technology that became very popular on the Web in the context of e-commerce. The most popular kind of link generation - dynamic recommendation of relevant links - was originally explored by InterBook and Hy-SOM (Kayama & Okamoto, 1999). ELM-ART pioneered another kind of link generation: generating links for similarity-based navigation. An example of link generation for similarity-based navigation in Knowledge Sea system (Brusilovsky & Rizzo, 2002) is shown on Figure 2.

ACT-R Lesson Units
Unit 1: Understanding Production Systems
Section 1.1: The ACT-R Production System
1.1.1 Declarative Units in ACT-R
1.1.2 Production Rules in ACT-R
1.1.3 Production Rule Format
1.1.4 ACT-R's Condition Form
1.1.5 ACT-R's Action Side

Content
Glossary
Help
Search
Interface

Back

Teach this Page

1.1.2 Production Rules in ACT-R

A production rule is a statement of a particular contingency that controls behavior. Examples might be

```
IF the goal is to classify a person
    and he is unmarried
THEN classify him as a bachelor

IF the goal is to add two digits d1 and d2 in a column
    and d1+ d2 = d3
THEN set as a subgoal to write d3 in the column
```

The condition of a production rule (the IF part) consists of a specification of a goal and a number of chunks while the action of a production rule (the THEN part) basically involves the creation or modifications of some chunks. The above is an informal English specification of production rules. You will learn the syntax for their precise specification within the ACT-R system.

A production rule specifies an [   ] to be taken when a [   ] is met.

Continue   Teach me

Background:

procedural knowledge

Outcome:

action
condition
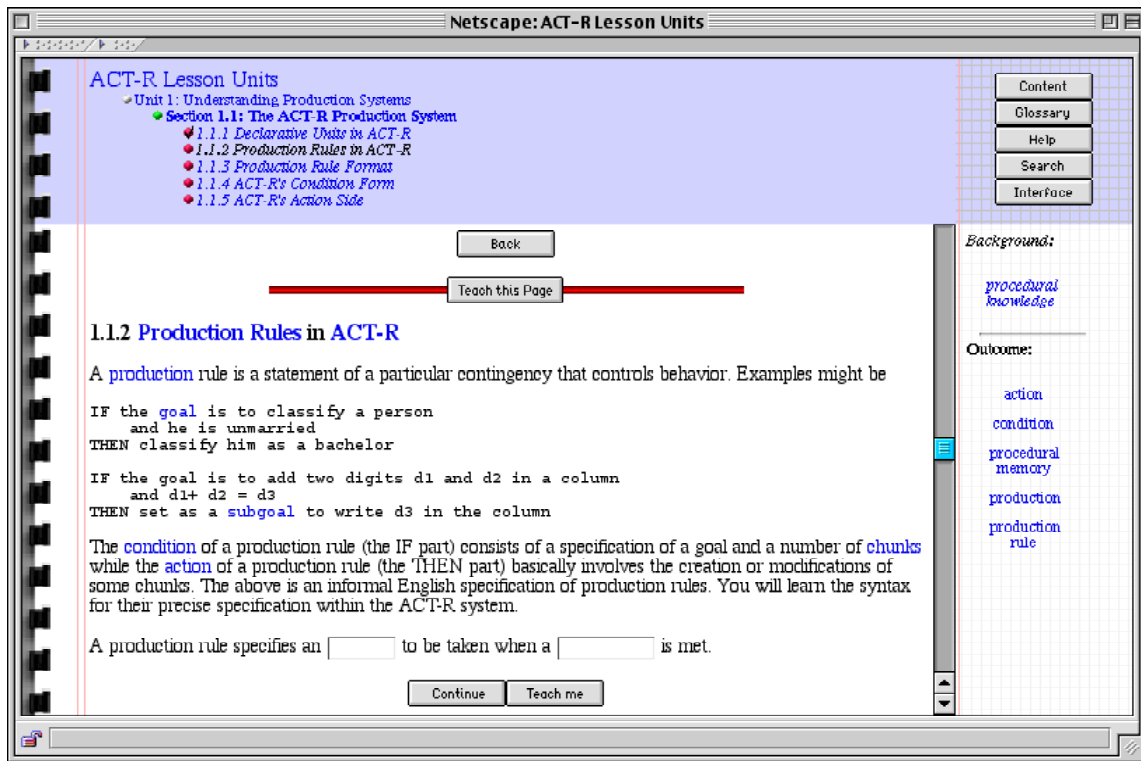procedural memory
production
production rule

Figure 1: Adaptive navigation support in InterBook. Adaptive annotation is provided in the form of colored bullets; adaptive guidance is accessible through "Teach Me" and "Teach this Page" buttons.

2.2 Early Studies of Adaptive Navigation Support

Early empirical studies of adaptive navigation support techniques were focused on the need to confirm the benefits of adaptation. These studies featured a "with or without" approach and focused on the use of audit trails to provide information about the effectiveness of link annotation in terms of improving browsing or learning efficiency. In the first published study, de La Passardiere & Dufresne (1992) conducted experiments with MANUEL EXCEL focussing on the value of history-based mechanisms such as a three-stage footprint (unseen, partially seen and completed) and the use of adaptive advice. Their work pointed to the value of an adaptive history-based mechanism as a means of navigation support.

In a year after that Kaplan et al. (1993) have reported another evidence in favor of adaptive navigation support. They performed two pilot studies with their system HYPERFLEX. In the first study they examined the usefulness of goal-directed search in the hypertext and demonstrated that goal-based adaptive sorting seriously decrease search time and the number of searched topics, while the correctness of answers even increased

slightly. In the second study they have compared the efficiency of two specific ways of adaptation in HYPERFLEX - interest-based and goal-based. Three versions of HYPERFLEX were used in experiment: the version with interest-based adaptation only, the version with goal-based adaptation only, and a fully functional system with both kinds of adaptation. The authors concluded that both methods of adaptation are efficacious, because the users of the fully functional system showed better performance then either of the two other groups. While the authors were not able to report about the significance of the results, they show that sorting-based adaptive navigation support can improve user performance in information search tasks.

A comprehensive study of hiding and annotation has been performed by Brusilovsky and Pesin (1998). The study compared student learning performance with three versions of ISIS-Tutor system: a non-adaptive version, a version with adaptive annotation, and a version with both adaptive annotation and hiding. Twenty-six computer science freshmen took part in the experiment. To complete the course, each user had to accomplish the same learning goal: to learn ten concepts and to solve ten problems. The subjects were divided randomly into three groups that used three versions of the system. The results of the study have demonstrated several benefits of adaptive navigation support in reducing the navigation efforts to achieve the education goal. In particular, the overall number of navigation steps and the number of task repetitions (ie, trials to solve previously visited task) were significantly smaller for both versions with adaptive navigation support. The average time to achieve the goal has been also reduced dramatically in both adaptive versions, however this difference was not significant due to the large individual variance.

2.3 The InterBook Study: The First Challenge

The results of both HYPERFLEX and ISIS-Tutor studies of adaptive navigation support were very encouraging a coherent. The results have demonstrated that three different adaptive navigation support techniques - sorting, hiding and annotation - are efficient adaptation techniques. These techniques can improve user performance in hypermedia by significantly reducing navigation difficulty. With these kinds of adaptive navigation support, the user can achieve the same result generally faster and with using a significantly smaller number of navigation steps. Adaptive annotation and hiding in an educational context can reduce user's floundering in the hyperspace and make learning with hypermedia more goal-oriented.

With this encouraging results in hands, we have attempted a similar study to evaluate another implementation of adaptive navigation support by link annotation used in InterBook system (Brusilovsky et al., 1998). The study was originally reported in (Brusilovsky & Eklund, 1998). Here we consider some of the results of this study from a different prospect, in the context of our current understanding.

The study involved 25 undergraduate teacher education students in an educational computing elective at the University of Technology, Sydney. The students were exposed to two chapters of a textbook about *ClarisWorks* databases and spreadsheets, and used the *InterBook* system both *with* and *without* adaptive link annotation (the version without adaptive annotation had no checkmarks and all bullets were green regardless of the link status). The goal of this experiment was to assess what impact, if any, user model-based link annotation would have on students' learning and on their paths through the learning space.

The experiment took place over a four-week period. In the first two-hour session, students were introduced to *InterBook* and its features were explained to them. They used the system for an hour, and answered a questionnaire about its features. This questionnaire showed that almost all students were familiar with what each of the buttons and annotations meant. They were then free to use the system at any time during the following week. In the second session, students were randomly divided into two groups of equal size, one group receiving link annotation, while the other group did not. They were allowed access to the chapter of the textbook on databases, which had been authored into InterBook, and they completed a questionnaire. Students had access to the database chapter for the following week. In the third session, students took a multiple-choice test on the database section of the textbook. InterBook navigation logs were analyzed along with the test results and the questionnaire responses.

It was a real surprise for our research group to discover that there *no significant results* between the two groups in regards of student final knowledge of the subject. A two-sample T-test showed that there was no significant difference at the 0.05 level in the test means for those with ANS and those without ANS. While students seem to understand and like adaptive navigation support (ANS) features, it didn't influence their performance on tests. It didn't took us long to find the explanation of this result. The analysis of student navigation traces have shown clearly that about 80% of student navigation steps were made either with Continue and Back buttons (that were tools for sequential navigation in InterBook) or with hot words in text. Neither these two buttons, nor the hot words were annotated in the experimental version of *InterBook*. The students essentially were not using the adaptive annotation of non-sequential links. It is hardly surprising that ANS has provided no significant difference. While some more detailed analysis has demonstrated a number of benefits of adaptive annotations (especially for those students who were using it), the overall results of the experiment were clear: adaptive annotation is not a silver bullet and may not work in some contexts.

We should admit that at that time, blindfolded by the success of earlier experiments, we made a wrong conclusion. Instead of thinking what could be wrong with the adaptive annotation, we decided to blame the experiment design - ie, missing annotations of "next" and "previous" buttons - and to repeat the experiment with some

small modifications (such as having all links properly annotated). Needless to say that our new experiment with a similar group of subjects hasn't brought any significant results either (Eklund & Sinclair, 2000). Despite of all links now being annotated, the dominated majority of students still preferred sequential navigation and got no benefits from multiple adaptively annotated links offered by InterBook at each page of the course.

2.4 The Role of Student Knowledge Level

Our current interpretation of the results of our InterBook study is different. We think that adaptive annotation was simply an improper adaptation technology for the category of students who took part in both InterBook experiments. Indeed, teacher education students in their majority had neither knowledge of ClarisWorks database, nor any experience that could be relevant to this subject. It is known from the educational hypertext research that this kind of "total novice" hypertext users tend to follow a sequential way of navigation, and ignore links that can get them out of the linear path. In this context, adaptive link annotation that helps the students to choose the most relevant non-sequential links has little chances to be useful. In fact, the results of the InterBook experiment can be considered as a good success since the experiment has shown also that the students who received adaptive annotation were much more eager to explore non-sequential links and that adaptive annotation benefited those who did use it. In contrast, the subjects of the ISIS-Tutor experiment were computer science students with some good background knowledge and a solid experience of work with similar systems. These students were eagerly using provided non-sequential links and benefited a lot from adaptive annotation.

Thus, the two studies together demonstrate clearly that adaptive annotation is not a silver bullet, as we tend to think originally. It can be helpful for some categories of users and relatively useless for others. A few other empirical studies performed by our colleagues from the original ELM-ART team provides some further evidence that student knowledge of the subject as well as their hypertext experience is a critical parameter that determines success of popular adaptive navigation support techniques.

The first of the studies performed with ELM-ART system was reported in (Weber & Specht, 1997). The study attempted to evaluate the role of two kinds of adaptive navigation support - adaptive annotation and adaptive direct guidance - in helping the users to work with ELM-ART system. The critical parameter was the user knowledge of the subject (LISP and programming) as evaluated by a self-assessment pre-test. All users were simply visitors of ELM-ART Web system. Since nobody forced the users to work with the system, it was assumed that the longer a user stays with the system, the more supportive was the used version of the system for him or her.

All visitors were randomly assigned to one of the four different treatments. The first pair of treatments contrasted the adaptive annotation of links with usual annotation

performed by WWW browsers that annotate links that have already been visited and cached. The second pair of treatments contrasted adaptive direct guidance with "next" button with a version without this button. The results of this study showed that subjects who had no previous experience with any programming language stayed longer with ELM-ART when they were guided by the system using a NEXT button. In turn, most subjects who were already familiar with at least one other programming language (and also were familiar with Web browsers) were not affected by adaptive guidance and stayed with the learning system longer when links were annotated adaptively.

Similar results on the connection between starting level of knowledge and relevant adaptation techniques were presented in (Specht & Kobsa, 1999). This paper (that was also a result of re-processing and earlier study) has compared the value of adaptive annotation and adaptive hiding. It has found an interaction effects of the post-hoc variable "previous knowledge of learners" and the adaptive treatments. The interaction effects have an impact on learners' scores in knowledge tests, the time learners needed to browse adaptive hypertexts, the number of their page requests, and the type of information requested by them. More specifically, the paper provided a good evidence that learners with higher previous knowledge seem to prefer non-restricting adaptive methods, wile learners with low previous knowledge can profit from the guidance of more restrictive adaptive methods.

2.5 The Knowledge Sea Study

While the objective role of various navigation support techniques (time spent, resulting level of knowledge, eagerness to work with the system) were explored extensively, we have not found any studies that evaluate the user subjective opinion about these techniques. In our recent study of goal-based guidance and link generation in the Knowledge Sea system (Brusilovsky & Rizzo, 2002) we have asked students to rate different aspects of the system and attempted to compare these ratings for the students with different levels of knowledge of the subject. Knowledge Sea applies a neural network-based mechanism to process a large number of pages from different Web-based tutorials along with a set of closed corpus documents (such as lecture notes) and group then by similarity. As a result, a user with a specific educational goal - such as to do readings associated with a particular lecture - can use an automatically generated list of relevant links to explore (Figure 2).
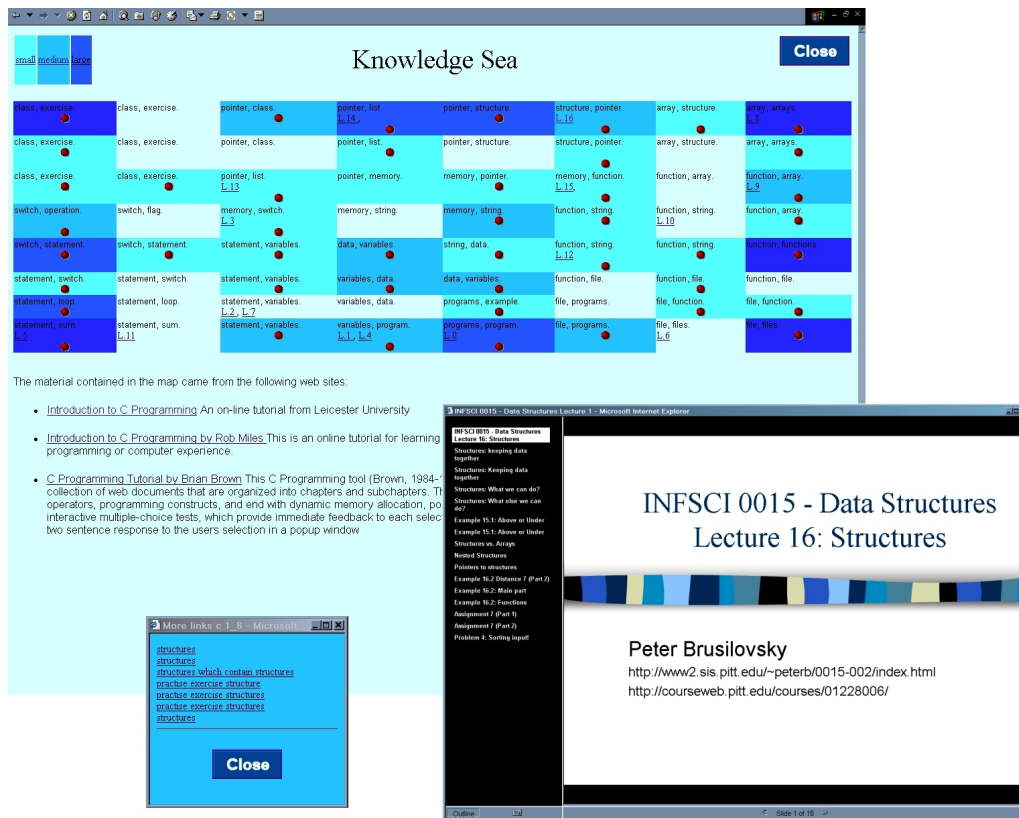
Figure 2: Link generation and similarity-based navigation in Knowledge Sea system.

The subjects in our Knowledge Sea study were students of an undergraduate course on Programming and Data Structures at the University of Pittsburgh. The study was formally announced to the students about 2 week before the final exam. Knowledge Sea was introduced to the students 2 week before the final exam as a tool that can help to find additional material on the topics involved into the final exam and thus achieve a better level of understanding. The system was available to all 39 students of the class. Right before the final exam, the students who used the tool for at least 20-30 minutes were asked to fill in a on-line questionnaire to collect their opinion about different aspects of the system. Twenty-one students choose to participate.

Of 13 total questions presented in the questionnaire, 10 are relevant to the focus of this paper. The questions were designed in a Likert 4-point style where answer 1 always corresponded to a very positive opinion, answer 2 to a positive, answer 3 to neutral or light negative, and answer 4 to negative. (ie, positive answers are labeled with lower numbers). The questions were grouped in 5 classes. The questions A and B are aimed at understanding the general judgment (overall feeling) about the system. Question C regards the judgment about the idea of an information map as a tool that collects external

resources and lesson handouts. The questions D and E are related to the interface of the system. Question F is related to the quality of the open corpus tutorial pages used to build the system. The effectiveness of the clustering and information organization is assessed using the questions G, H, I.

To explore the correlation between the student's knowledge of the subject and their opinion about Knowledge Sea, we have decided to consider student's final course grade as a measure of their knowledge of the subject. We have divided the students into three groups – A-level students (that includes A+, A, and A-), B-level students (B+, B, and B-) and C-D level students. To start with, just from the pattern of student participation we can notice clear that the better is the grade, the more excited were the students about the system. 8 of 11 A-level students, 10 of 17 B-level students and only 3 of 11 students with lower grades choose to take the questionnaire. The same trend is reflected by the profile of student answers. (Figure 5). For all questions the average opinion of the A-level students was quite more positive than the opinion of B-level students. It is also almost always more positive than the opinion of C-level students, but we have decided to exclude the C-D group from consideration since this group includes only 3 students. Note that the student positive opinion was not influenced by their grade. The study was completed before the students took the final exam. Our hypothesis is that the system was simply more helpful for the students with better knowledge of the subject.

Overall, 62% to 95% of all students (depending on the question) evaluated different features of the system positively or very positively. Yet, the students with good knowledge of the subject were happier about all the system aspects. A large difference between A-level and B-level student satisfaction in the system's interface hints us that it is the interface of the system that created most problems to the students with weaker knowledge of the subject. Indeed, a careful analysis of the student free form comments hints has indicated that the students with the lower grades were often confused by the large number of links generated and organized by the system and have troubles in selecting the most relevant link to follow. This is quite consistent with the experiments summarized above as well as with some other studies of educational hypermedia. Once more, we are receiving evidence that offering a relatively large number of navigation opportunities, even organized and annotated adaptively, for the students with lower knowledge is not a wise idea though it works well in the same context for the student with a higher level of knowledge.

One interesting fact to notice is that there is one question where the opinion of the students with different grades is about the same. This is the question B that asked the students to isolate a bit from their personal feelings and answer how relevant will be the system for the class as a whole. As we see, the students were quite consistent in thinking "for the whole class". The feedback of B-level students (and C-D-level students too, while it is not shown) was a bit more positive than their feedback on similar questions A

and C and, vice versa, the opinion of A-level students was a bit less positive. It was looking like students themselves were able to understand that the system was really more suitable for the students with better knowledge of the subject.
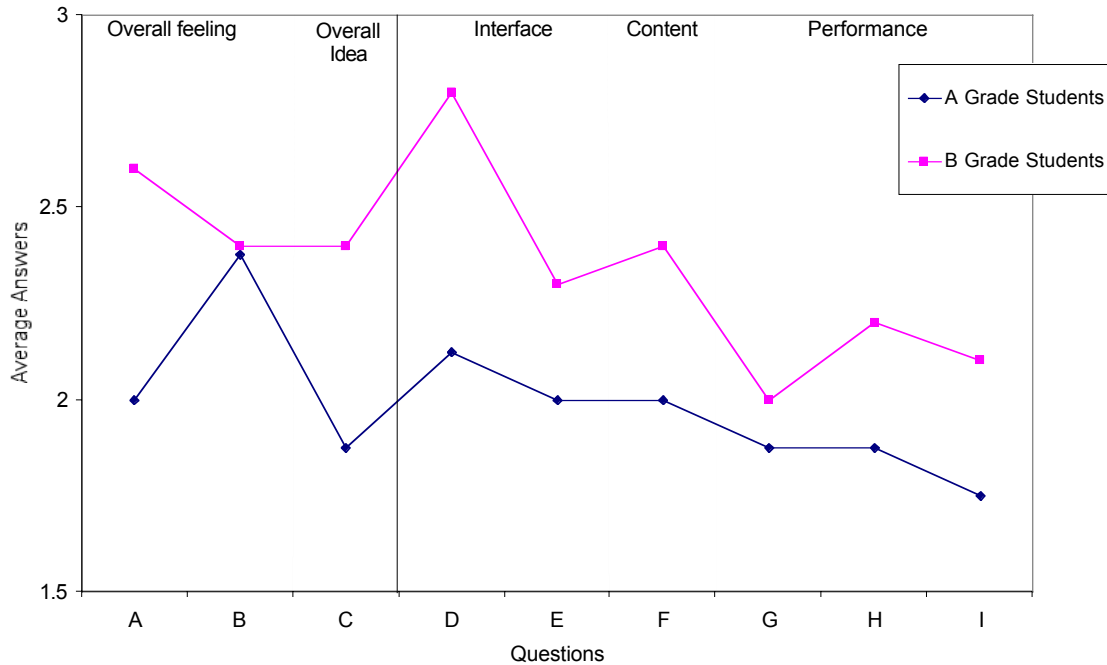

Figure 3: Student rating of the Knowledge Sea aspects split by student knowledge of the subject as evaluated by the final course grade.

3. Discussion: The Need for Meta Adaptation

The results of several empirical studies of adaptive navigation support techniques calls adaptive hypermedia researchers to reconsider the traditional perception of this popular approach to adaptation. The positive results of early studies generated unrealistically high expectations of adaptive navigation support. Even nowadays, many researchers consider adaptive navigation support techniques reviewed at the beginning of this paper as a kind of silver bullet. Once a known adaptive navigation support technique is implemented in a system, it will magically work benefiting to all users working with it

At the same, the researchers who are carefully evaluating their systems have already several opportunities to learn that each of known adaptive navigation support technology has some applicability limits. Depending on the context, an adaptive navigation support technology may or may now work properly. Moreover, different known techniques may work most efficiently in different context.

This paper has attempted to review known adaptive navigation support techniques and examine their behavior along one dimension of possible differences between

application context: the level of user knowledge of the subject. The studies that we have summarized in the paper have provided very good evidence that users with different knowledge level of the subject may appreciate different adaptive navigation support technologies. It seems that for users with little or no knowledge of the subject, most relevant are restrictive technologies such as direct guidance or hiding that guide them by adaptively limiting their navigation choice. In contrast, most relevant technologies for the users with some reasonable knowledge level of the subject are "rich" linking technologies such as adaptive annotation and multiple link generation.

This paper calls for two clear conclusions. First, we should conclude that researchers and practitioners in the field of adaptive hypermedia should be more careful with selecting adaptive navigation support technologies for their target systems. Instead of using an ad-hoc approach selecting the technology that looks nicer or that is easier to implement, they have to examine existing empirical evidence and make the selection appropriate for the context. Second, it should be observed that more careful studies of known technologies in different are required on the way to making adaptive navigation support an everyday tool. The existing studies just can't provide sufficient guidance in selecting the proper techniques. In brief: we have too many developed techniques and too few studies.

Both these two conclusions are rather obvious and may not worth writing a separate paper. There is, however, the third conclusion that was one of the main motivations behind this paper: the need for meta-adaptation. Until now, the focus of research in adaptive hypermedia was selecting or developing an adaptation technology. The technology then is expected to take care about adapting to an individual user. What we are learning now is that the very selection of the adaptation technology should be adaptive to the class of users at hand. The natural step forward is to develop meta-adaptive hypermedia systems. A meta-adaptive system should have a number of different adaptation technologies at its disposal. It should also be aware about the limits of applicability of every technology and be able to adaptively select the very adaptation technology that fits best the given user and the given context. It is also natural to expect that meta-adaptive systems will be able to constantly extend their own knowledge about the applicability of different technologies by observing the success of these technologies in different context and learning from these observations.

The emergence of meta-adaptive systems seems to be a natural process of evolution of adaptive systems. A similar process can be observed in the neighboring area of intelligent tutoring systems. While classic works in this field were focusing on developing and evaluating different teaching strategies, following research has shown that different strategies work best in different context. More lately work on teaching strategies was focused on meta-adaptation, ie, on formally representing different teaching strategies and developing a mechanism (Van Marcke & Vedelaar, 1995). Most recent

stream of work is related with recording data about the effectiveness of different teaching strategies and learning from these data.

The author hopes that this paper provides both good food for though for the users and developers of adaptive navigation support technologies ands a good case for developing meta-adaptive hypermedia systems.

References

Brusilovsky P (1996) Methods and techniques of adaptive hypermedia *User Modeling and User-Adapted Interaction* 6 2-3, 87-129.

Brusilovsky P (2001) Adaptive hypermedia *User Modeling and User Adapted Interaction* 11 1/2, 87-110, also available at http://www.wkap.nl/oasis.htm/270983.

Brusilovsky P and Eklund J (1998) A study of user-model based link annotation in educational hypermedia *Journal of Universal Computer Science* 4 4, 429-448, also available at http://www.iicm.edu/jucs_4_4/a_study_of_user.

Brusilovsky P, Eklund J, and Schwarz E (1998) Web-based education for all: A tool for developing adaptive courseware *Computer Networks and ISDN Systems* 30 1-7, 291-300.

Brusilovsky P and Pesin L (1998) Adaptive navigation support in educational hypermedia: An evaluation of the ISIS-Tutor *Journal of Computing and Information Technology* 6 1, 27-38.

Brusilovsky P and Rizzo R (2002) Map-based horizontal navigation in educational hypertext *Journal of Digital Information* 3 1, also available at http://jodi.ecs.soton.ac.uk/Articles/v03/i01/Brusilovsky/.

Brusilovsky P, Schwarz E, and Weber G (1996) ELM-ART: An intelligent tutoring system on World Wide Web in Frasson C, Gauthier G and Lesgold A (eds.) *Intelligent Tutoring Systems*. Lecture Notes in Computer Science, Vol. 1086, Springer Verlag, Berlin, 261-269, also available at http://www.contrib.andrew.cmu.edu/~plb/ITS96.html.

De Bra P M E (1996) Teaching Hypertext and Hypermedia through the Web *Journal of Universal Computer Science* 2 12, 797-804, also available at http://www.iicm.edu/jucs_2_12/teaching_hypertext_and_hypermedia.

de La Passardiere B and Dufresne A (1992) Adaptive navigational tools for educational hypermedia in Tomek I (ed.) *Computer Assisted Learning*. Springer-Verlag, Berlin, 555-567.

Eklund J and Sinclair K (2000) An empirical appraisal of the effectiveness of adaptive interfaces for instructional systems *Educational Technology and Society* 3 4, also available at http://ifets.ieee.org/periodical/vol_4_2000/eklund.html.

Kaplan C, Fenwick J, and Chen J (1993) Adaptive hypertext navigation based on user goals and context *User Modeling and User-Adapted Interaction* 3 3, 193-220.

Kayama M and Okamoto T (1999) Hy-SOM: The semantic map framework applied on an example case of navigation in Cumming G, Okamoto T and Gomez L (eds.) *Advanced Research in Computers and Communications in Education*. Frontiers ub Artificial Intelligence and Applications, Vol. 2, IOS Press, Amsterdam, 252-259.

Specht M and Kobsa A (1999) Interaction of domain expertise and interface design in adaptive educational hypermedia in Brusilovsky P and Bra P D (eds) Proceedings of Second Workshop on Adaptive Systems and User Modeling on the World Wide Web, Toronto and Banff, Canada, 89-93.

Van Marcke K and Vedelaar H (1995) Learner adaptivity in generic instructional strategies in Greer J (ed) Proceedings of AI-ED'95, 7th World Conference on Artificial Intelligence in Education. Washington, DC, 323-333.

Weber G and Brusilovsky P (2001) ELM-ART: An adaptive versatile system for Web-based instruction *International Journal of Artificial Intelligence in Education* 12 4, 351-384, also available at http://cbl.leeds.ac.uk/ijaied/abstracts/Vol_12/weber.html.

Weber G and Specht M (1997) User modeling and adaptive navigation support in WWW-based tutoring systems in Jameson A, Paris C and Tasso C (eds.) *User Modeling*. SpringerWienNewYork, Wien, 289-300, also available at http://www.psychologie.uni-trier.de:8000/projects/ELM/Papers/UM97-WEBER.html.