The Design of the TERENCE Adaptive Learning System

Tania Di Mascio[1], Rosella Gennari[2], Pierpaolo Vittorini[1]
1:University of l'Aquila, Italy, tania.dimascio@univaq.it,pierpaolo.vittorini@cc.univaq.it
2: Free University of Bozen-Bolzano, Italy, gennari@inf.unibz.it

Abstract: Nowadays, circa 10% of 7-11 olds turn out to be poor comprehenders: they demonstrate text comprehension difficulties, related to inference-making, despite proficiency in low-level cognitive skills like word reading. The use of more intelligent ALSs to custom-tailor reading interventions in the form of smart games for poor comprehenders has tremendous potential. The TERENCE ALS embodies that potential. This paper presents the design choices of the TERENCE ALS. First it specifies the preliminary user requirements of the system, and the system's intended usage through selected usage scenarios. The paper continues describing the system's learning material, made of stories and games, according to the preliminary users' requirement specifications. The learning material and stories are structured as in the presented domain model, part of the conceptual model of the TERENCE architecture. Finally, the paper explains the architecture of the TERENCE system, focusing on the adaptive engine and intelligent backbone.

1 Introduction

Developing the capabilities of children to comprehend written texts is key to their development as young adults. Text comprehension skills and strategies develop enormously from the age of 7-8 until the age of 11, when children develop as independent readers. Henceforth, we refer to 7-8 to 11 olds as novice (text) comprehenders. Nowadays, more and more novice comprehenders turn out to be poor (text) comprehenders: they demonstrate difficulties in deep text comprehension, despite well developed low-level cognitive skills like vocabulary knowledge, e.g., see [Cain et al., 2001] for hearing poor comprehenders, and [Marschark et al., In press] for deaf poor comprehenders. The following introductory subsections briefly outline:

1. the preliminary requirements of poor comprehenders in relation to text comprehension, according to state of the art, and hence the motivations and rationale for the design of our TERENCE adaptive learning system (ALS) for poor comprehenders,
2. a background on ALSs and an overview of related work.

The paper then continues with usage scenarios that illustrates the usage of our ALS. It then shows the story, game and student models of the conceptual model of the ALS, and ends by describing the system's logical architecture.

1.1 Preliminary User Requirements and Rationale of Our Work

In the United Kingdom, poor comprehenders are estimated to be 5% to 10% of novice comprehenders [Lyon et al., 2003]. Similar numbers were identified in Italy [Cornoldi et al., 1996], as well as in other countries [Ehrlich and Rémond, 1997,Nation and Snowling, 1998,Paris et al., 2005]. The estimate dramatically increases when the whole population of hearing-impaired children is considered. For instance, in [van Bon et al., 2006] the authors estimate that only 19% out of 504 hearing impaired 7-20 olds have reading comprehension scores above the third grade level. Several studies experimentally demonstrate that poor comprehenders fail to master the following reasoning skills in processing written stories, e.g., see [Cain et al., 2001]:

(s1) coherent use of cohesive devices such as temporal connectives,

(s2) inference-making from different or distant parts of a text, integrating them coherently,

(s3) detection of inconsistencies in texts.

Nowadays, there is clear evidence that such reasoning skills (s1, s2, and s3) are very likely to be causally implicated in the development of deep text comprehension. In particular, experiments show that inference-making questions centred around (s1, s2, and s3), together with adequate visual aids, are pedagogically effective in fostering deep comprehension of stories, e.g., see [National Reading Panel, 2000].

However, finding stories and educational material that are appropriate for poor comprehenders is a challenge, and hence educators are left alone in their daily interaction with poor comprehenders. Most reading material for novice comprehenders is paper based, and is not easily customisable to the specific requirements of poor comprehenders, e.g., in the types, number or position of temporal connectives. Few ALSs promote general reading interventions, but they have high-school or university textbooks as reading material, instead of stories, and are developed for old children or adults, and not specifically for poor comprehenders.

Our TERENCE ALS aims at filling such a gap: its reading material (in English and in Italian languages) will be stories adapted to the specific requirements of poor comprehenders, and its reading interventions will be interactive question-games centred around reasoning skills, like (s1), (s2), and (s3) above, that foster the development of deep text comprehension, both accompanied by adequate visual aids.

1.3 Background on ALSs and Related Work

The conceptual model of an ALS is usually made up of (see [Santos et al., 2003]) (i) the student model that describes the student's main features, the (ii) domain model that structures the learning material, the (iii) environment model for the hardware/software capabilities, and the (iv) adaptation model that, given the previous models, characterises the actual adaptive mechanism and adaptation engine. In order for ALSs to be pedagogically effective, the adaptation components usually require more specialised AI-techniques, tailored to the specific needs of their users [Brusilovsky, 2000,Paramythis and Loidl-Reisinger, 2004,Shute and Zapata-Rivera, 2007], than those of a classical intelligent tutoring system. Both the planner approach of the MIGRAINE system [Buchanan et al., 1995] and the progress-based guidance mechanism of QuizGuide [Brusilovsky et al., 2004] are particularly relevant to the adaptation engine of our ALS. The MIGRAINE system uses a knowledge base and a text to generate questions to patients that are trying to comprehend a medical report on their condition, an approach that leads to a better understanding and a greater satisfaction. QuizGuide, instead, proposes quizzes to students. Quizzes are grouped in sets, according to their topics. QuizGuide attempts to guide students to the most relevant self-assessment quiz sets, by tracing correct and incorrect answers for the quizzes, calculating mastery levels for each quiz. These levels are propagated to the topic level, forming the mastery view of the whole topic. The evaluation of progress-based navigation support in QuizGuide demonstrated that this technology has succeeded in guiding towards the most appropriate quiz sets.

As for the student model, different aspects of user modelling were studied independently from different angles. Aside from distribution, scalability and performance aspects [Kobsa and Fink, 2006] and context information [Jameson, 2001], the principal motivations for the development of user models are (i) to characterise an individual user and (ii) have a generic representation of types of users. The former has received most attention research and proof-of-concept implementations. For instance, the KBS-Hyperbook [Henze and Nejdl, 1999] and TRAILS projects [Heller et al., 2007] based their modelling on (reasoning over) logged user actions. In the AHA! project [De Bra et al., 2006], on the other hand, user actions are typically not logged but immediately translated into higher-level user model information, and Brusilovsky and Millán [Brusilovsky and Millán, 2007] focus on Bayesian networks. As Kay [Kay, 2008] already points out, with maturing ALSs, long-term usage, and interoperability issues among extant systems, the second option becomes more important. There are, however, few ontologies described in the literature, primarily being the generic user models of GUMO [Heckmann et al., 2005] and GRAPPLE [De Bra et al., 2010]. The fact that they are generic allows us to reuse them and make them more specific for the domain and students of our ALS. To the best of our knowledge, there are no user or domain models that, nowadays, cover adequately our specific subject domain of poor comprehenders, and deep text comprehension skills.

2 Usage Scenarios

1) A learner, Carol, interacts with TERENCE by herself.

Carol is an 8 year old girl. She is smart, but utterly alone. Her parents work all day and have little time for reading stories and interacting with Carol. Carol prefers playing with video games that are interactive, and interests her more than reading paper stories.

Carols wants to try reading the TERENCE stories by herself, at her own pace. She knows TERENCE because she frequently uses it in class with her teacher.

Carol opens TERENCE that shows her different books: one book of fantasy stories, one of dinosaurs, etc. She chooses the fantasy book and the Layla Princess as avatar. She loves this avatar because Layla has a wonderful diadem.

TERENCE opens a spatial map with several lands. Carol chooses the story of the Magic Castle. TERENCE shows the story and Carol starts reading. Alas, some of the sentences are too long and obscure to her. She asks TERENCE for stories written in a simplified language. TERENCE, through its adaptive engine, provides her with a simplified story of the fantasy book. She starts reading again; this story is really funny and easy to read!

Carol, who is enthusiastic about the characters of the story, asks TERENCE for playing games. TERENCE shows her smart games concerning the story, with an intuitive visual interface. The games and the visual interface are funny, intuitive, and coherent with the fantasy book's environment and the Layla Princess avatar. Carol starts playing and obtains a good score. When she decides to quit TERENCE, she is surprised in seeing that own avatar (the Layla Princess) has gained a new diadem better then the first one! She decide to read more tomorrow – she wants Layla to have an entire parure!

2) An educator, Alice, browses TERENCE and sets TERENCE for a specific poor comprehender.

Alice teaches in a primary school in Italy. Giorgio is one of her learners. He is 10 year old and he loves vampires. He reads quickly, but seems unable to understand what he reads. The school psychologist administered him the Neale Analysis of Reading Ability test (NARA, e.g., see the Revised British Edition, Neale, 1997). The administration of the NARA requires the child to read a set of stories aloud, and to answer a number of questions after each story. Giorgio read all the words correctly, but was found unable to answer the comprehension questions.

Alice wants to use TERENCE because she heard that TERENCE may help her with Giorgio due to its visual games. Alice doesn't know what adaptive learning systems are, and uses the computer only for browsing the Internet or writing documents. Despite this, Alice opens TERENCE and starts browsing the educator GUI. She reads the help for educators, and chooses some gothic stories. She tries playing with the associated smart games, and she discovers the possibility to set TERENCE for Giorgio's characteristics, like linguistic skills and interests, by means of a user friendly interface. She is now sure that TERENCE can be of aid to Giorgio.

Alice sets the characteristics of Giorgio, e.g., linguistic skills and interests, and she asks TERENCE for adequate stories for Giorgio. TERENCE shows her a list of adequate gothic book stories, adequate according to the system. Alice approves of the system's choice.

Giorgio starts reading and is immediately enthralled by the gothic environment where he has chosen the avatar of the "Conte Dracula" trying to win against the "Werewolf".

3) An educator, Joshua, uses TERENCE to choose stories and smart games, and to monitor his own learners.

Joshua is the teacher of a 4th year class in Sussex. In his class, there are learners with deep text comprehension problems. It is Christmas time and, using TERENCE, he reads a Christmas story to the class. Alex, Rose, Evelyn and Joe are poor comprehends. They cannot sum up the main episodes of the story, and ask Joshua for clarifications.

Joshua decides not to explain them the story, and instead encourages them to play with some smart games. First of all, Joshua uses TERENCE to select some specific smart games for each of his learners. Since Joshua had stored the characteristics of the learners in TERENCE, the TERENCE system can select the appropriate games concerning the Christmas story, for each learner.

Alex, Rose, Evelyn and Joe try solving the funny Christmas games!

Alex chooses as avatar Santa Claus: when Alex scores well on the games, Santa Claus hands out chocolates!

After the smart game session, Joshua evaluates the evolution of his learners, using TERENCE.

Joshua notes that Alex, Rose, Evelyn and Joe read the entire story, and, even more surprisingly for him, all of them could tackle the majority of the games in which they had to reconstruct the temporal flow of the story, and summarise its main episodes.

Satisfied by this, Joshua hands out chocolates to his learners, feeling a bit like a novel Santa Claus.

4) An educator, Mary, uses TERENCE to create stories and smart games for a specific poor comprehender.

Mary is the teacher of Ted, a 11 year old boy, with serious text comprehension difficulties. He gets easily distracted and is tired of reading after three or four sentences. However, Ted loves solving riddles, logic puzzles and spends hours in playing chess on the Internet.

Mary knows TERENCE and thinks it can improve Ted's reading skills on account of its smart games. Mary sets the characteristics of Ted in TERENCE and then lets Ted choose a book. Ted selects the chess book. Then Mary selects some specific events of the book's stories, and asks TERENCE to build smart games with them for reasoning about the stories chosen by Ted. Mary asks for "easy" games; TERENCE selects the games that challenge Ted to find inconsistencies in stories, after he reads them.

The smart games are presented to Ted by the King Avatar of chess, the avatar that Ted has chosen for his book. The smart games' interface is funny, intuitive, and coherent with the chess environment. Moreover, Mary can provide input on demand to the adaptive engine of TERENCE, and hence help fine-tune the guidance mechanism for Ted.

According to TERENCE, Ted scores well on the most difficult games of the set! Therefore Mary asks TERENCE for other sets of smart games for the same stories, which are more demanding according to TERENCE. Then Mary continues following the performance of Ted online.

5) An expert, Steve, who is a natural language processing researcher, uses TERENCE to monitor its automatic annotation of stories.

Steve is a researcher at the University of Rome and responsible of the natural language processor of TERENCE . Steve uploaded a new book of stories in TERENCE, and now wants to check the XML annotations of those stories.

The new book of TERENCE is about traditional fables that Steve loves very much. When Steve reads the annotation of a story, he notes that the not all the TLINK annotations are present in the "Little Red Riding Hood". In particular, a "before" TLINK between two consecutive events is clearly missing. Thus he asks TERENCE to introduce it, and TERENCE adds the required "before" TLINK between the selected events.

6) An expert, Katrina, who is a psycho-linguist, uses TERENCE to work on the stories of TERENCE.

Katrina is a psychologist expert of poor text comprehenders, part of the TERENCE team that is responsible for studying the parameters for classifying stories in difficulty classes for poor comprehenders. She opens TERENCE to check the classification of the new stories of the book about witches and witchcraft, "The Bewitched Lake" book. She notes that a story of the book is automatically classified by TERENCE as "easy". However, in her opinion, several sentences of this story are demanding for poor comprehenders, and hence this story should be classified as "average". She then asks TERENCE to revise the parameters for the classification of the story. After her revision, TERENCE classifies the story as "average".

3 The Conceptual Model of the TERENCE ALS

The preliminary analysis of the context of use and the user requirements already allows us to characterise both the learning material and goals of our ALS, and hence draft the conceptual model, as outlined in the following subsections.

3.1 The Learning Material

The main learning material of our ALS consists in illustrated stories, and interactive smart games centred around inference-making questions for reasoning about the stories.

Stories. Even if finding stories for novice comprehenders is not difficult, finding stories that are adequate to poor comprehenders and challenging is difficult. For instance, the comprehension of causal-temporal relations in text is problematic for poor comprehenders. However, the use of appropriately interspersed causal-temporal connectives in stories facilitate their comprehension of the stories (see Section 1). Therefore, TERENCE evaluates the appeal and adequacy of its stories with different experts of poor comprehenders (e.g., Katrina in above sixth scenario) and types of poor comprehenders, e.g., with different ages, grammar skills and, in case of deaf children, different first-language.

Smart games. The reasoning smart games are based on a story's explanatory events. Preliminary prototypes of the games are being designed by following assessed interventions for stimulating the reasoning skills on texts outlined in the motivations. For instance, according to our preliminary user requirements' analysis, poor comprehender are in need of games that pose and solicit questions about specific features of events in the stories, monitoring the students' comprehension of the story flow. Let us make two concrete examples of such question-games, starting from the following excerpt of a famous children's story.

Mummy Duck watches the big egg but sees no signs of cracking. So she decides to keep sitting on it. After some days, while she is sitting on it, an ugly grey duckling cracks the big eggshell.

The games of our ALS will feature who-questions like

(q1) Who sits on the big eggshell?

or questions that are possibly more difficult like

(q2) Does the big eggshell crack before Mummy Duck watches it?

for reasoning about *implicit* causal-temporal information in the story. Such questions will be accompanied by adequate visual aids that help children foster their deep comprehension of the related information, in line with the results of [Arfé et al., 2009].

3.2 The Main Submodels of the Conceptual Model

The main submodels of the conceptual model of our ALS is specified as follows:

(s1) the learning material described above is structured in the domain model, which is composed of two main submodels: one for stories, the other for smart reasoning games

(s2) the student model for the characteristics of poor comprehenders, divided in emotional and cognitive ones;

(s3) the adaptation model for the teaching process, specifying the rules correlating concepts of the domain model and the student model.

The models are being developed as ontologies starting from existing ones. In particular, the current story ontology, part of the domain ontology, analyses and correlate features of texts that are difficult for poor comprehenders. The ontology was constructed with a middle-out approach: after analysing the reading comprehension taxonomies overviewed in [Chikalanga, 1992] and the Coh-metrix indices in light of our preliminary user requirements, we created the concepts of our story ontology. For instance, the story ontology has as "difficult word" as concept, specialised in non-frequent word, abstract word, unfamiliar word. It has also different concepts of inference-making about specific features of events of stories, e.g., the actor of an event-action, the attributes of the actor.

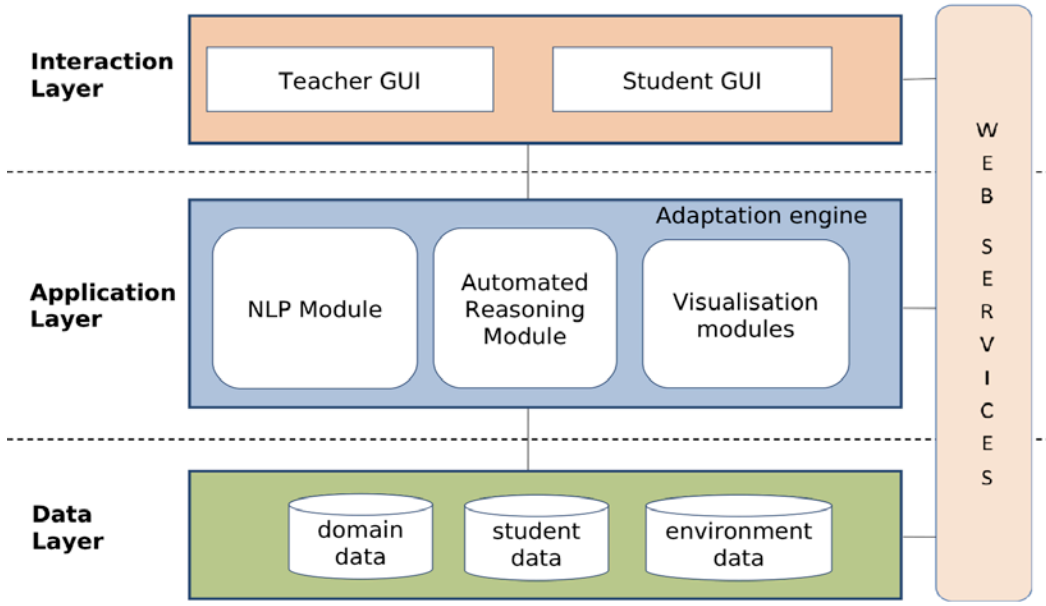

Figure 1. Architecture and adaptation engine

4 The Architecture of Our ALS

The current logical architecture of our ALS is divided into three layers, namely: the data layer, the application layer, and the interaction layer. See Figure 1. The data layer stores the stories and smart games (according to the domain model), the information specific to each learner (according to the student model), and the information about the supported clients (according to the environment model). The application layer implements the adaptation engine, and the intelligent backbone responsible for the automated feedback on games (e.g., correct or incorrect answers to games). Finally, the interaction layer contains the users' interfaces.

The following two subsections detail the intelligent backbone of the application layer of the system, and the role of web services in it.

4.1 The Intelligent Backbone of the Adaption Engine

The main modules of the adaptation engine of our ALS are a constraint-based automated reasoner and a natural language processing (NLP) module, consisting of a processor for English stories and one for Italian stories, that lay at the core of the engine and constitute its intelligent backbone, so to speak. A similar idea was proposed in the LODE web e-tool for deaf children for the first time in [Gennari and Mich, 2007]. In TERENCE, the work is carried several steps ahead as follows: the NLP modules serve to annotate stories with specific XML tags; the tags are used for classifying stories as in the story ontology, for generating and answering smart games as well.

The way in which the adaption engine assesses the performance of a student on games can be manifold. In particular, the approaches of MIGRAINE [Buchanan et al., 1995] and QuizGuide [Brusilovsky et al., 2004] seem the most relevant for TERENCE, in that they use questions for assessing the users' understanding of a specific domain as TERENCE does.

According to the assessment of the students' performance on a class of games, our TERENCE ALS will attempt to guide the students to the most adequate games, following the game ontology of the domain model, and to the most adequate story, following the story ontology of the domain model.

4.2 The Role of Web Services

Adaptive hypermedia techniques have proven their ability to support an improved user learning experience, albeit with some exceptions. Nevertheless, there has been no widespread adoption of these techniques so far [Brusilovsky and Henze, 2007]. One of the reasons is that ALSs are usually monolithic, and they lack of re-usability and interoperability. Our TERENCE ALS will advance on this by virtue of its architecture based on ontologies and web services, as well as by using automated annotation and reasoning tools. In the specific context of our ALS, tasks will be implemented through web services and their composition. For instance, let us refer to the generation of a temporal question-game like question (q2). The composition will go as follows:

1.firstly, a NLP web service operation, which takes the story as input and returns the annotations in the novel annotation language, is invoked;

2.secondly, these annotations are taken as input, and an operation of the Automated Reasoner service deduces further causal-temporal annotations as output, updating them;

3.finally, the updated annotations are taken as input and a further operation of the NLP web service generates as output the example question-game.

A first clear advantage of such an approach is that it allows for the reuse of our architecture in other languages, by implementing the appropriate NLP services. Furthermore, since web services are accessible through HTTP calls, they can be invoked directly in their respective organisations. Finally, the high-level operations might also be implemented with a programming-in-the-large paradigm, e.g. BPEL [Weerawarana et al., 2005], thus allowing for an easy deploy of further operations, that become web services, and therefore re-usable to build up more complex tasks.

5 Conclusions

In this paper, we motivated the need of an ALS for poor comprehenders, with a first analysis of their preliminary user requirements, explained also by means of several usage scenarios. Then the paper outlined the current conceptual model of our ALS's architecture, focusing on its adaptive engine. The intelligent backbone of this mainly consists of an NLP module and a constraint-based automated reasoner for generating question-games for reasoning about stories, and providing an automated feedback to the games.

Currently, we are working on a finer-grained analysis of the preliminary user requirements, to be reviewed through evaluations; the updated requirements will provide the input for the revision of the conceptual model of TERENCE outlined in this paper. Such evaluations, more in general, will serve to assess the usability of our ALS and in particular: (1) the appeal and adequacy of its learning material, (2) the pedagogical effectiveness of our ALS in improving the text comprehension of poor comprehenders.

Acknowledgements. The research of the three authors of this paper was supported by the European

Community's Seventh Framework Programme FP7/2007-2013 under the TERENCE grant agreement n. 257410. The second author's work was also supported by a grant of the CARITRO bank of Trento for the LODE project.

References

Arfé, B., Gennari, R., and Mich, O. (2009). Evaluations of the LODE Temporal Reasoning Tool with Hearing and Deaf Children. In Cognitive and Metacognitive Educational Systems (MCES2009), AAAI Fall Symposium.

Baker, F. (2001). The Basics of Item Response Theory. ERIC Clearinghouse on Assessment and Evaluation, University of Maryland, College Park, MD.

Brusilovsky, P. (2000). Adaptive hypermedia: From intelligent tutoring systems to web-based education. In Lecture Notes In Computer Science,, 5th international Conference on intelligent Tutoring Systems.

Brusilovsky, P. and Henze, N. (2007). Open corpus adaptive educational hypermedia. In The Adaptive Web. Methods and Strategies of Web Personalization, Lecture Notes in Computer Science 4321, pages 671–696.

Brusilovsky, P. and Millán, E. (2007). User models for adaptive hypermedia and adaptive educational systems. In The Adaptive Web. Methods and Strategies of Web Personalization, Lecture Notes in Computer Science 4321, pages 3–53.

Brusilovsky, P., Sosnovsky, S., and Shcherbinina, O. (2004). Quizguide: Increasing the educational value of individualized self-assessment quizzes with adaptive navigation support. In Nall, J. and Robson, R., editors, World Conference on E-Learning, pages 1806–1813. AACE.

Buchanan, B. G., Moore, J. D., Forsythe, D. E., Carenini, G., Ohlsson, S., and Banks, G. (1995). An intelligent interactive system for delivering individualized information to patients. Artificial Intelligence in Medicine, 7(2):117–154.

Cain, K. (In Press). Young Readers' Processing and Comprehension of Connectives.

Cain, K., Oakhill, J. V., Barnes, M. A., and Bryant, P. E. (2001). Comprehension Skill, Inference Making Ability and their Relation to Knowledge. Memory and Cognition, 29:850–859.

Chikalanga, I. (1992). A Suggested Taxonomy of Inferences for the Reading Teacher. Reading in a Foreign Language, 8.

Cornoldi, C., De Beni, R., and Pazzaglia, F. (1996). Profiles of reading comprehension difficulties: An analysis of single cases, pages 113–136. Reading comprehension difficulties: Processes and interventions. Lawrence Erlbaum Associates.

De Bra, P., Smits, D., and Stash, N. (2006). The Design of AHA! In ACM Conference on Hypertext and Hypermedia.

Di Mascio, T., Gennari, R., Lang, I., and Vittorini, P. (2010). Visual Tools for Annotating Temporal Expressions with TimeML: a Critical Overview. KRDB Technical Report.

Duran et al., 2006. Duran, N., McCarthy, P., Graesser, A., and McNamara, D. (2006). Using Coh-Metrix Temporal Indices to Predict Psychological Measures of Time. In Conference of the Cognitive Science Society, pages 190–195. Cognitive Science Society.

Ehrlich, M. F. and Remond, M. (1997). Skilled and less skilled comprehenders: French children's processing of anaphoric devices in written texts. British Journal of Developmental Psychology, 15:291–309.

Gennari, R. and Mich, O. (2007). Constraint-based Temporal Reasoning for E-learning with LODE. In Proc. of CP 2007. Springer.

Goldin-Meadow, S., and Mayberry, R (2001). How Do Profoundly Deaf Children Learn to Read? Learning Disabilities Research and Practice, 16:221–228.

Graesser, A.C., McNamara, D., Louwerse, M.M., and Cai, Z.. Coh-metrix: Analysis of Text on Cohesion and Language. Behavior Research Methods, Instruments, & Computers, 36(2):193–202, 2002.

Heckmann, D., Schwartz, T., Brandherm, B., and Kreoner, A. (2005). Decentralized User Modeling with

UserML and GUMO. In Workshop on Decentralized Agent Based and Social Approaches to User Modelling (DASUM 2005), pages 61–65.

Heller, J., Levene, M., Keenoy, K., Albert, D., and Hockemeyer, A. (2007). Cognitive aspects of trails: A stochastic model linking Navigation Behaviour to the Learner's Cognitive State Navigation Behaviour to the Learner's Cognitive State. Trails in Education. Technologies that Support Navigational Learning. Sense Publishers.

Henze, N.and Nejdl, W.(1999).Adaptivity in the KBS Hyperbook System. In 2nd Workshop on Adaptive Systems and User Modeling on the WWW, workshop held in conjunction with the World Wide Web Conference (WWW8) and the International Conference on User Modeling.

Jameson, A. (2001). Modeling both the Context and the User. Personal and Ubiquitous Computing, 5:29–33.

Kay, J. (2008). Lifelong Learner Modeling for Lifelong Personalized Pervasive Learning. IEEE Transactions On Learning Technologies, 1(4).

Kobsa, A. and Fink, J. (2006). An LDAP-based User Modeling Server and its Evaluation. User Modeling and User-Adapted Interaction, 16(2):129–169.

Lyon, G., Fletcher, J. M., and Barnes, M. (2003). Learning Disabilities. Child Psychopathology. The Guilford Press.

Marschark et al., In press. Marschark, M., Sapere, P., Convertino, C., Mayer, C., W. L., and Sarchet, T. Are Deaf Students' Reading Challenges Really About Reading?

Musselman, 2000. Musselman, C. (2000). How Do Children Who Can't Hear Learn an Alphabetic Script? A Review of the Literature on Reading and Deafness. Journal of Deaf Studies and Deaf Education, 5:9–31.

Nation and Snowling, 1998. Nation, K. and Snowling, M. J. (1998). Individual Differences in Contextual Facilitation: Evidence from Dyslexia and Poor Reading Comprehension. Child Development, 69:996–1011.

National Reading Panel (2000).Teaching Children to Read: An evidence-based Assessment of the Scientific Research Literature on Reading and its Implications for Reading Instruction. Technical report, National Institute of Child Health and Human Development.

Paramythis, A. and Loidl-Reisinger, S. (2004). Adaptive Learning Environments and e-Learning Standards. Electronic Journal of e-Learning, 2(1):181–194.

Paris, S. . G., Carpenter, R. D., Paris, A. H., and Hamilton, E.E. (2005). Spurious and Genuine Correlate of Children's Reading Comprehension, pages 131–160. Children's Reading Comprehension and Assessment. LEA.

Santos, J.,Anido, L., Llamas, M.,Á lvarez,L., and Mikic, F. (2003). Computational Science, chapter Applying Computational Science Techniques to Support Adaptive Learning, pages 1079–1087. LNCS, Springer.

Shute, V. J. and Zapata-Rivera, D. (2007). Handbook of Research on Educational Communications and Technology (3rd Edition), chapter Adaptive technologies. Taylor & Francis.

van Bon, L. N., Wauters, W. H. J., and Tellings, A. E. J. M. (2006). Reading Comprehension of Dutch Deaf Children. Reading and Writing, 19:49–76.

Weerawarana, S., Curbera, F., Leymann, F., Storey, T., and Ferguson, D. (2005). Web Services Platform Architecture: SOAP, WSDL, WS-Policy, WS-Addressing, WS-BPEL, WS-Reliable Messaging, and More. Prentice Hall PTR.

Zwaan, R. (1999). Narrative Comprehension, Causality, and Coherence: Essays in Honor of Tom Trabasso, chapter Five Dimensions of Narrative Comprehension: the Event-indexing Model. Erlbaum.