Indirectly Visible Bayesian Student Models

Diego Zapata-Rivera
Research and Development Division
Educational Testing Service
Princeton, NJ 08536

Abstract

Inspectable Bayesian student models have been used to support student reflection, knowledge awareness and communication among teacher, students and parents. This paper presents a new approach to interacting with inspectable Bayesian student models called indirectly visible Bayesian student models. In this approach, the student model is seen through the eyes of a pedagogical agent (e.g., a virtual student). This approach has been implemented in the context of an Assessment-Based Learning Environment for English grammar (English ABLE), where the student is asked to help a pedagogical agent find grammatical errors on various sentences. Since the pedagogical agent's knowledge levels, which are also the student's knowledge levels, are always visible, the student can see how much the pedagogical agent "knows" based on his/her actions. Initial reactions to this approach have been positive. We are planning on integrating it into assessment-based learning and gaming environments as indicators of progress that continuously change in light of new evidence.

1 INTRODUCTION

Assessment information can be obtained from a variety of sources including standardized assessments, classroom quizzes, group activities, and self- or negotiated assessment activities. Intelligent Tutoring Systems (ITSs) continuously monitor student performance and adapt their behavior to a changing view of the student maintained by the system (i.e., a student model).

Student models generally maintain rich student assessment information. Assessment information, when shared with students, teachers and parents, can be used to support formative dialogue in the classroom that can promote student learning. Black and Wiliam (1998a, 1998b), for example, established a clear link between formative assessments (assessment for learning) and classroom learning.

Open student models (OSMs) consider teachers, students, and sometimes parents to be more than just consumers of assessment information. In OSM, these participants play an active role by observing, updating, and acting based upon student model assessment information. OSMs have been used to support student reflection, knowledge awareness, group formation, student model accuracy and learning (Brna et al., 1999; Bull & Pain, 1995; Hartley & Mitrovic, 2002, Kay, 1998; Dimitrova, 2004; Zapata-Rivera & Greer, 2004).

Inspectable, interactive Bayesian student models have been used to integrate various sources of evidence (e.g., the system's and the student's view of the student model). Several visualization techniques including animation have been used to show how evidence of student performance is added to and propagated throughout the Bayesian student model (Zapata-Rivera & Greer, 2001, 2004).

Although various representational and interaction techniques have been used to implement OSMs, students always see the student model as the system's view of his/her knowledge, skills and abilities. This direct approach to OSMs confronts the learner with a view of the student model that could (or could not) match that of his/her own requiring the student to react to it. Students could react in a variety of ways depending on many factors including student self-esteem, personality traits, and personal beliefs regarding computers in general. For example, while some students could respond in a negatively way categorically rejecting the system's claims leaving no room for negotiation, some could, instead, try to understand the system's claims in detail and perhaps even challenge them, some would just accept them, and some would completely ignore them without even looking at them.

What if the system refers to a third person instead, for example, someone the student wants to help? Could such an approach avoid or at least attenuate some of these possible negative reactions? How would students react to this approach? We have implemented an indirect approach to interacting with Bayesian student models that capitalize on the idea of learning by teaching. In this approach students "teach" a pedagogical agent by providing help finding grammatical errors. Students can see whether the pedagogical agent is making progress (or not) by looking at how the indirectly visible student model changes and how the pedagogical agent reacts.

The indirectly visible Bayesian student modeling approach has been implemented in the context of an Assessment Based Learning Environment for English grammar called English ABLE. English ABLE makes use of a Bayesian student model that is used by pedagogical agents to provide adaptive feedback and adaptive sequencing of tasks. A view of the Bayesian student model is presented to the student through the eyes of a pedagogical agent.

This paper describes the Bayesian student model used in English ABLE, explains how the indirectly visible student model was implemented, describes its potential to be integrated into existing games, reports on initial student reactions, and concludes by discussing some open research issues and plans for future work.

2 ENGLISH ABLE

English ABLE is an Assessment-Based Learning Environment for English grammar. Assessment-based learning environments make use of assessment information to guide instruction.

English ABLE demonstrates the reuse of existing high-stakes tasks in lower stakes learning contexts. English ABLE currently draws upon a database of TOEFL® Computer-Based Testing (CBT) tasks to create new packages of enhanced tasks targeted towards particular component ELL skills.

In English ABLE, students try to help a virtual student (Carmen or Jorge) learn English by correcting this student's writing from a notebook of facts (sentences —enhanced TOEFL® tasks). Supplemental educational materials about specific grammatical structures are offered by a virtual tutor (Dr. Grammar).

Figure 1 shows a screenshot of English ABLE. The student is helping Jorge find grammatical errors within several sentences. The student selects an option and clicks on "*Check Answer*." Dr. Grammar offers verification feedback "*I see you have selected 'created'. However, this part of the sentence is correct.*," and additional adaptive instructional feedback (i.e., rules, procedures, examples and definitions). Students can ask Dr. Grammar for hints "*Ask for a hint*" before committing to a particular choice. In that case, Dr. Grammar provides a general rule related to the current grammatical structure. Students can also type a possible correction "*Suggested word*." Both asking for help and providing corrections are treated differently when adding evidence of student performance to the Bayesian model. Jorge's knowledge levels, which are also the student's knowledge levels (*indirectly visible Bayesian student model*), show a lack of knowledge for agreement. Jorge seems confused and expresses it "*I don't understand how to make the verb agree with the rest of the sentence.*"

Knowledge levels representing the pedagogical student's knowledge of English grammar are taken directly from the Bayesian network that supports the system (i.e., Bayesian student model). Although only three knowledge bars are shown in Figure 1 (i.e., *Agreement*, *Wrong Form* and *Omission/Inclusion*), a detailed view of the Bayesian student model containing information about low-level concepts is available upon student request (Details button).

2.1 Bayesian Student Model

Several authors in different areas have explored the use of Bayesian belief networks to represent student models (Collins et al. 1996, Conati et al. 2002, Horvitz et al. 1998; Mislevy & Gitomer, 1996; VanLehn & Martin, 1997; Reye, 2004).

English grammar can be divided into three main categories: use, form, and meaning (Celce-Murcia & Larsen-Freeman, 1999). We worked with experts to elicit an initial Bayesian structure for a student model (see Figure 2). The current structure of the Bayesian student model deals with English grammar form, although it could be extended to cover use and meaning.

Three sentence-level grammatical categories (i.e., *Agreement*, *FormofWord or Wrong Form*, and *OmissionInclusion*) have been chosen based upon a difficulty analysis that was performed using student data from native Spanish speakers. These three sentence-level grammatical categories are further divided into low-level sub-categories (leaf nodes) according to parts of speech (e.g., agreement has been divided into 3 leaf nodes: noun agreement, verb agreement, and pronoun agreement). Leaf-nodes are linked to 2 main knowledge areas (i.e., individual parts of speech: noun, verb and pronoun, and sentence-level grammatical categories).

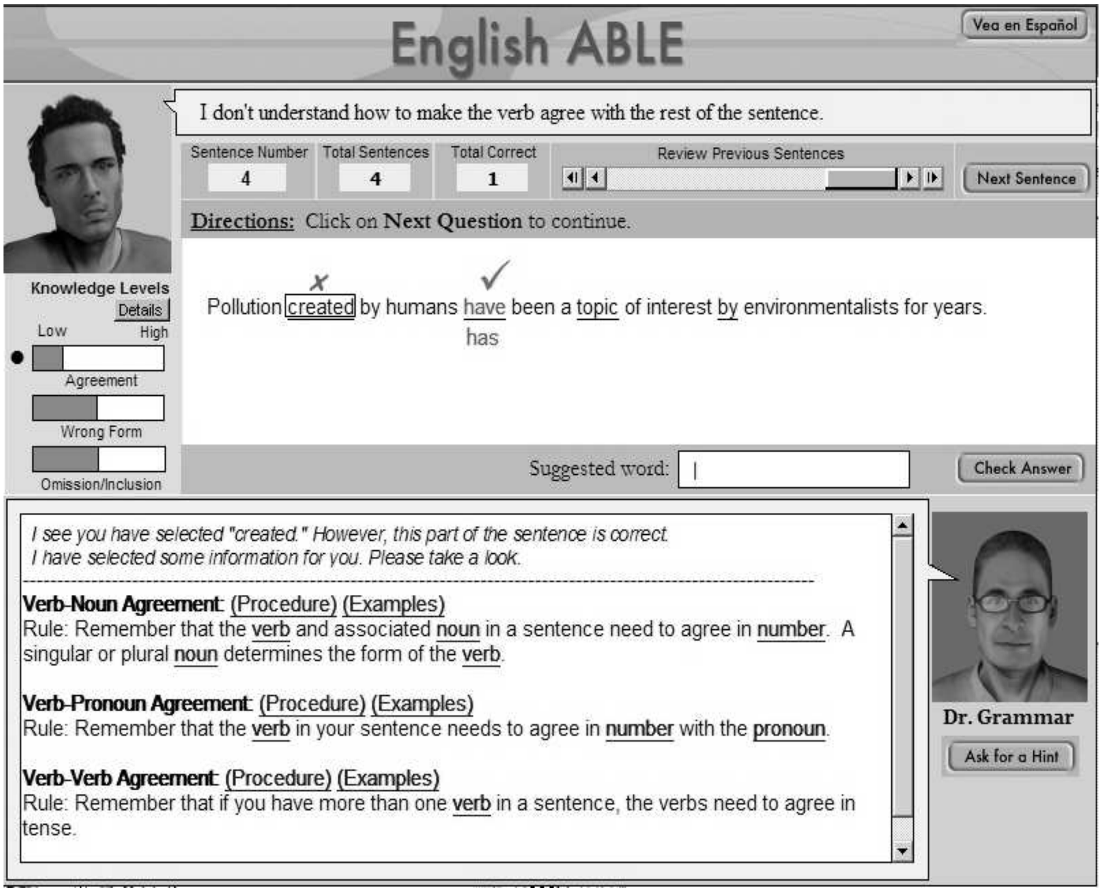

Figure 1: English ABLE.

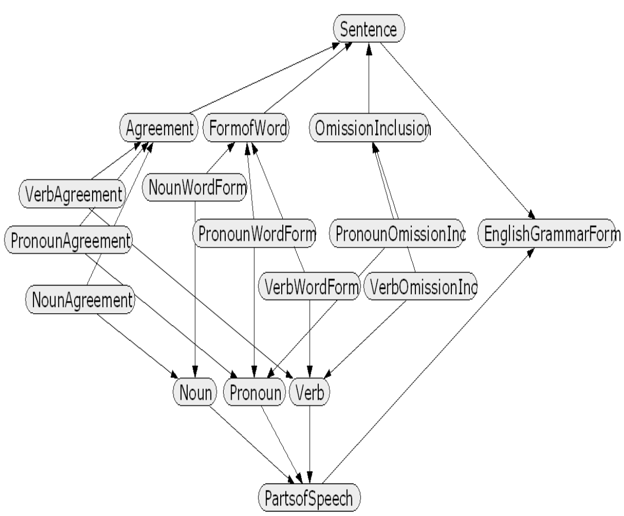

Figure 2: Bayesian student model.

Preliminary difficulty analysis plus data from experts were used to generate prior and conditional probabilities for the latent structure. Experts used a qualitative inspired method to produce probability values based on estimates of the strength of the relationship between any two variables in the model (Daniel, Zapata-Rivera & McCalla, 2003).

Each task was attached to a single category using existing classification metadata and corresponding Item Response Theory (IRT) discrimination and difficulty parameters (Lord & Novick, 1968; Embretson & Reise, 2000). Tasks were recalibrated (i.e., new IRT parameters were computed) based on data from all native Spanish speakers who took the test. The IRT-2PL model is described by the following formula:

$$\Pr(task_i = correct \mid prof_j) = \frac{1}{1+(e^{-1.701*a(prof_j-b)})},$$

where b is the difficulty parameter ($-3 \leq b \leq +3$, typical values for b), a is the discrimination parameter ($-2.80 \leq a \leq +2.80$, typical values for a), and $Prof_j$ represents an ability level (continuous proficiency variables were discretized using the following ability values: $Advanced = 0.96$, $IntermediateAdvanced = 0$, and $Intermediate$ = -0.96. These values come from quantiles of a normal distribution (Almond, et al. 2001)).

Figure 3 shows how tasks are connected to leaf-nodes using IRT parameters. Table 1 shows the resulting conditional probability table of Task 2 (a=1.5, and b= 0.4).

As the student makes progress (i.e., answers additional tasks), more tasks are dynamically added to the model. Observed values per task (i.e., correct or incorrect) provide evidence (as defined by its conditional probability table) to update the student model. Asking for help (*"Ask for a hint"*) and providing corrections (*"Suggested word"*) are handled by slightly adjusting the difficulty level of the task.

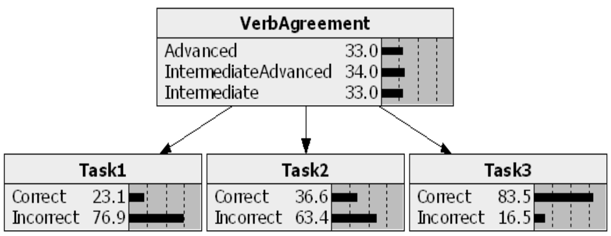

Figure 3: Three tasks connected to a leaf-node.

Table 1: Conditional probability table for Task 2 (a=1.5, b= 0.4)

	Pr(Task2|VerbAgreem)	
VerbAgreement	**Correct**	**Incorrect**
Advanced	0.807	0.193
IntermediateAdvanced	0.265	0.735
Intermediate	0.030	0.970

This underlying Bayesian network supports the knowledge levels and the pedagogical agents' behavior. That is, indirect knowledge levels are computed based on the corresponding probability distribution of a particular node. Pedagogical agents query the Bayesian student model to implement adaptive algorithms (i.e., adaptive feedback, adaptive sequencing of items, and adaptive behavior).

This Bayesian student model can be made available to students using a variety of approaches. For example, we could have used ViSMod (Zapata-Rivera & Greer, 2003) to show students a complete view of the graphical structure using visualization techniques such as node color, link size and animation to represent marginal and conditional probabilities. Although presenting the whole the Bayesian network can help students understand how the Bayesian student model works (e.g., understanding integration and propagation of evidence), it requires students to spend some time understanding and interacting with the student model. Interactive, collaborative and negotiated approaches to open student model use the student model as a communication tool engaging students in a formative dialogue aimed at supporting metacongition.

We do not have to show the whole Bayesian network to provide students with a sense of progress (e.g., weak and strong areas). We can just show an overall view covering main concepts/nodes or relevant ones depending on the tasks that the student is currently working on. Although, in this approach just a piece of the Bayesian student model would be open to students at a particular time, the whole internal Bayesian network is available to other components in the system (e.g., pedagogical agents). Different views of the Bayesian structure can be created to support the goals of the learning environment. These views can range from static student or teacher reports to interactive adaptive applications.

2.2 Pedagogical Agents

Pedagogical agents (e.g., Chan & Baskin, 1990; Graesser, Person, Harter, & TRG, 2001; Johnson, Rickel, & Lester, 2000) have been used to facilitate learning by supporting human-like interaction with computer-based systems. Pedagogical agents can act as virtual peers or virtual tutors. Pedagogical agents can model human emotions and use this information to facilitate learning (e.g., Picard, 1997; Nkambou et al., 2003).

An interesting variant of pedagogical agents are teachable agents (Biswas et al., 2001), which have been used to facilitate student learning. The student's role in these environments is to teach an artificial student how to act in a simulated environment. Students in English ABLE are asked to help a pedagogical agent (i.e., Carmen and Jorge) find grammar errors. Carmen and Jorge "learn" based on the student's performance. Students can see how much the pedagogical agent knows about a particular concept by looking at the indirectly visible Bayesian student model and by observing Carmen's and Jorge's changes in emotional states and associated utterances (Zapata-Rivera et al., 2007). Figure 4 depicts Jorge, Carmen and Dr. Grammar.

2.3 Indirectly Visible Bayesian Student Model

Bull et al. (2005) reported that children, university students and instructors understood and used a variety of student model external representations. However, they also warn of possible negative effects when low-performance students explore student models of more capable students (i.e., some of these students reported a negative effect on their motivation level and esteem).

Figure 4: Jorge, Carmen and Dr. Grammar.

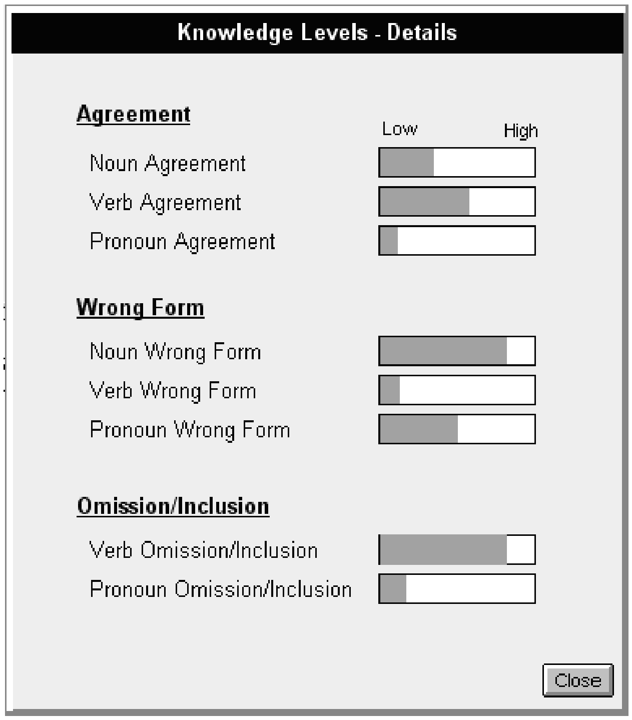

Figure 5: Jorge's knowledge levels.

English ABLE supports indirect inspection of Bayesian student models. We believe that exploring one's student model via a pedagogical agent is less intimidating and has the potential to foster student learning without the possible negative effects on self-esteem and motivation, especially for those students who are having a hard time with the system.

Previous research on inspecting Bayesian student models through the use of guiding artificial agents showed that agents can facilitate student interaction with the model by helping students navigate and find conflicting nodes. Guided agent interaction was linked to higher levels of student reflection (Zapata-Rivera & Greer, 2004).

Changes in marginal probability distributions can be depicted by showing a graphical indicator per each state of the node (e.g., three bars, one per each state of a proficiency node). This approach uses a great deal of screen space and requires users to have some familiarity with probability distributions to make sense of multiple changes occurring as more evidence becomes available and added to the Bayesian student model. Alternatively, we could choose one state (e.g., $\Pr(Proficiency_j = Advanced \mid evidence)$ and show just one bar. However, this approach, would not necessarily be sensitive to variations on marginal probability values occurring on the neglected states of the node.

In English ABLE, the length of each bar is calculated based on an Expected A Posteriori (EAP) score that takes into account the whole marginal probability distribution of a particular node, producing a value that ranges from zero to 1. This EAP-length score is computed using the following formula:

$$Length_i = \frac{\sum_j^n C_j \Pr(proficiency_i = state_j)}{n},$$

where C_j is a constant numerical value assigned to each state of a node based on its proficiency level (i.e., *Intermediate* $= 0$, *IntermediateAdvanced* $= 1$, and *Advanced* $= 2$) and n is the index of the highest proficiency state (e.g., $n = 2$, in this case).

Figure 5 shows a detailed view of Jorge's knowledge levels. This view of the student model appears when the student clicks on ("Details") (see Figure 1).

Tables 2 and 3 show how marginal probability and EAP values change based on the student's responses to a series of tasks. EAP values capture slight variations of marginal probabilities. The final effect is an indicator bar that continuously adjusts as new evidence is added to the model.

Table 2: Sequence of Probability and EAP-length values for a student solving *NounAgreement* tasks. Marginal probability values converge to the *Intermediate* state as EAP-length values get closer to zero

P(Int)	**P(IntAdv)**	**P(Adv)**	**EAP**	**Resp**
0.647	0.280	0.073	0.21	Cor
0.429	0.423	0.147	0.36	Inc
0.676	0.297	0.027	0.18	Inc
0.833	0.164	0.004	0.09	Cor
0.684	0.306	0.010	0.16	Inc
0.832	0.167	0.001	0.08	Cor
0.686	0.311	0.003	0.16	Inc
0.831	0.169	0.000	0.08	Inc
0.918	0.082	0.000	0.04	

Table 3: Sequence of Probability and EAP-length values for a student solving *NounWordForm* tasks. Marginal probability values converge to the *IntermediateAdvanced* state as EAP-length values get closer to 0.5

P(Int)	P(IntAdv)	P(Adv)	EAP	Resp
0.156	0.368	0.476	0.66	Cor
0.064	0.343	0.593	0.76	Cor
0.024	0.295	0.681	0.83	Inc
0.102	0.560	0.338	0.62	Cor
0.004	0.593	0.403	0.70	Cor
0.001	0.521	0.478	0.74	Inc
0.005	0.801	0.194	0.59	Cor
0.002	0.754	0.244	0.62	Inc
0.005	0.916	0.078	0.54	

We are currently experimenting with fading as a mechanism for forgetting about old pieces of evidence and assigning more weight to more recent evidence. Views of past data can be handled by using windows of various sizes that implement various fading policies. These views of the student model can be maintained and dynamically adjusted based on student performance. For example, pedagogical agents and other consumers of student model information can maintain their own view into the past based on how important evidence of past performance is to accomplish their student learning goals.

Pedagogical agents (e.g., virtual tutors) implementing various forms of adaptive instruction use their own view of the student model to keep track of students progress. Some of these pedagogical agents can implement some form of collaborative or negotiated assessment using a view of the student model to support formative dialogue between students and teachers. Evidence gathered from these educational stakeholders can then be integrated with existing evidence of student performance into an aggregate view of the student model that implements a particular policy for integration of evidence. This framework can be used as a research testbed for studying the effects of several adaptive instructional and assessment strategies on student learning.

2.4 Indirectly Visible Bayesian Student Models and Games

Indirectly visible Bayesian student models can be integrated as part of first person role-playing games. In these games, each player chooses a character that identifies him/herself in the game. Each character has a particular personality, skills, and abilities. Some of these traits change during the game as the player makes progress in the game. Up-to-date estimates of players' competencies based on a Bayesian student model can be integrated into the game as progress/state indicators. Using these indicators, players see how their competencies are changing based on their performance in the game. This level of self-awareness can be linked to the development of meta-cognitive abilities.

We are planning to use embedded assessments to capture valued information without disrupting the flow and engagement of the game. We have started applying some of these ideas in the context of a popular first person role-playing game called The Elder Scrolls® IV:*Oblivion*™© (Bethesda Softworks, 2006). For more information about how indirectly visible Bayesian student models can potentially be integrated into existing games, see Shute, Ventura, Bauer & Zapata-Rivera (in press).

3 INITIAL STUDENT REACTIONS

We recently conducted a study focusing on usability issues and learning effects in relation to English ABLE tools and interface. We report on the results from our usability study. Information regarding learning effects can be found in Zapata-Rivera et al. (2007).

Participants included 149 native Spanish speakers (ESL students) who were assigned to 3 different conditions (i.e., test preparation, English ABLE simple and English ABLE enhanced). Forty six of the participants were assigned to the enhanced version of English ABLE that included: a Bayesian student model, an indirectly visible student model and pedagogical agents.

In general, we were interested in knowing how students reacted to the indirectly visible Bayesian student modeling approach. In particular, we wanted to know how students reacted to the pedagogical agents and their knowledge levels. Students were asked to respond to a series of questions using a likert scale with the following choices: strongly agree, agree, disagree, and strongly disagree.

Results from the usability study showed that 88% of the participants assigned to English ABLE enhanced, understood the knowledge levels presented in the indirectly visible student model, 86% thought that the knowledge levels were useful, and 86% agreed that the knowledge levels helped them understand what Jorge/Carmen knew.

Participants agreed with the following statements: (a) "I liked helping Carmen/Jorge find grammar errors" (90%), (b) "Carmen's/Jorge's comments were useful" (78%) , (c) "Helping Carmen/Jorge motivated me to keep going" (90%), (d) "I have helped Carmen/Jorge a lot by finding the grammar errors" (73%), (e) "I have learned by helping the Carmen/Jorge with his/her sentences" (90%), (f) "The feedback provided by Dr. Grammar helped me learn" (87%), and (g) "I think Carmen and Jorge liked my help" (81%).

In addition, some of the students' comments seemed to indicate that they understood their role as teachers and used student model information to continuously assess learning progress. For example, a motivated student mentioned that *"My Carmen is happy. Her knowledge levels are increasing,"* while a struggling student exclaimed: *"Poor Carmen, she is not learning a lot from me."*

Initial results show that students enjoyed the current implementation of the indirectly visible Bayesian modeling approach. We believe that teaching someone else and seeing how he/she makes progress (or not) can be a strong motivational factor that can help maintain students engaged in the learning process. Although initial results are encouraging more studies are needed.

4 DISCUSSION & FUTURE WORK

Different external representations can be used to offer views of the student model and interaction techniques can be implemented to help students and teachers to interact with the student model. It is important to take into account the goals of the learning session and the need of having an accurate student model in order to decide which kind of support is more appropriate for a particular situation (Zapata-Rivera & Greer, 2004).

Although students seemed to enjoy helping pedagogical agents find grammatical errors, current implementation of the agents was limited to providing additional scaffolding in a language accessible to students and showing various emotional states based on the current state of Bayesian student model. Interaction with these pedagogical agents could be enhanced by supporting dialogue based interaction. For example, pedagogical agents could ask students to explain particular actions or elicit additional information from students aiming at mapping the limits of their understanding regarding a particular topic.

Students could also question the estimates of knowledge assigned to the pedagogical agent. Does the agent really know about a particular grammatical structure? A student could think: *"Let's ask the agent some questions to see how he/she answer."* Testing the pedagogical agent on particular topics will also provide interesting evidence of student knowledge that can be added to the model. Should the pedagogical agent answer the questions at the level of the student or act as a weaker student? Should Dr. Grammar intervene if/when the student is teaching a wrong concept to the pedagogical agent or trying to game the system? How should agents respond to the questions raised by students? How do we convince students that their help is really helping the pedagogical agent "learn" the concepts? Although highly motivated students can engage in this kind of interaction with pedagogical agents, what kinds of mechanisms should be in place to maintain and encourage such high levels of motivation? How do we implement this level of interaction without negatively affecting the flow of a game? These are all interesting open research questions that motivate and inform our plans for future work.

Future work includes using assessment information to support learning in various contexts, harnessing the power of games and technology to provide highly interactive adaptive learning environments that seamlessly use assessment information to improve student learning, skills and performance in valued domain areas.

Acknowledgements

I would like to thank the members of the English ABLE team including: Malcolm Bauer, Thomas Florek, Waverly VanWinkle, Meg Powers, Debbie Pisacreta, Janet Stumper, James Purpura, Valerie Shute, Russell Almond, Jody Underwood, Margaret Redman, Christopher Pfister, Hae-Jin Kim, Linda Tyler, Cathrael Kazin, Jan Plante, Yong-Won Lee, Victor Aluise, Feng Yu, Mary Enright and Maurice Hauck. I also want to thank Valerie Shute, Russell Almond, Eric Hansen and three anonymous reviewers for providing valuable feedback on earlier versions of this paper. Also, I would like to extend my sincere appreciation to the students, teachers and school administrators who participated in our study.

References

R., Almond, L., Dibello, F., Jenkins, R., Mislevy, D., Senturk, L., Steinberg, & D., Yan (2001). Models for conditional probability tables in educational assessment. In T. Jaakkola and T. Richardson, editors, *Artificial Intelligence and Statistics*, 137-143.

Bethesda Softworks (2006) The Elder Scrolls® IV:*Oblivion*™© Bethesda Softworks LLC, a ZeniMax Media company. All Rights Reserved.

G., Biswas, D., Schwartz, J., Bransford, & the Teachable Agent Group at Vanderbilt (TAG-V) (2001). Technology support for complex problem solving: From SAD environments to AI. In K. D. Forbus and P. J. Feltovich (Eds.), *Smart machines in education: The coming revolution in educational technology.* Menlo Park, CA: AAAI/MIT Press. 71-97.

P., Black, & D., Wiliam (1998a). Assessment and classroom learning. Assessment in Education: Principles, Policy, and Practice, 5 (1), 7-74.

P., Black, & D., Wiliam (1998b). Inside the black box: Raising standards through classroom assessment. London: School of Education, King's College.

P., Brna, J., Self, S., Bull, & H., Pain (1999). Negotiated Collaborative Assessment through Collaborative Student Modelling.*Proceedings of the workshop Open, Interactive, and other Overt Approaches to Learner Modelling at AIED99.* Le Mans, France, pp. 35-44.

S., Bull, & H., Pain (1995). "'Did I say what I think I said, and do you agree with me?': Inspecting and Questioning the Student Model", *Proceedings of World Conference on Artificial Intelligence in Education (AACE)*, Charlottesville, VA, 501-508.

S., Bull, A. Mabbott, A.S. Abu Issa, & J. Marsh (2005). Reactions to Inspectable Learner Models: Seven Year Olds to University Students. *AIED'05 Workshop on Learner Modelling for Reflection, to Support Learner Control, Metacognition and Improved Communication between Teachers and Learners.* 23-32

M., Celce-Murcia & D., Larsen-Freeman (1999). The Grammar Book (2nd ed.). Heinle & Heinle Publishers.

T.W., Chan & A.B., Baskin (1990). Learning Companion Systems. In Frasson, C., and Gauthier, G. (eds.) *Intelligent Tutoring Systems: At the crossroads of AI and Education.* Ablex Pub., 6-33.

J. A., Collins, J. E., Greer, & S. X., Huang (1996). Adaptive Assessment using Granularity Hierarchies and Bayesian Nets. In Frason, C., C., Gauthier, G. and Lesgold, A., (eds.), *Proceedings of Intelligent Tutoring Systems ITS'96.* Berlin: Springer, 569-577.

C., Conati, A.S., Gertner, & K., VanLehn, (2002). Using Bayesian Networks to Manage Uncertainty in Student Modeling. *In User Modeling and User-Adapted Interaction.* 12, 371-417.

B., Daniel, J.D., Zapata-Rivera, & G., McCalla (2003). A Bayesian Computational Model of Social Capital in Virtual Communities. *Proceedings of the First International Conference on Communities and Technologies; C&T 2003.* Kluwer Academic Publishers. ISBN 1-4020-1611-5.

V., Dimitrova (2003). StyLE-OLM: Interactive Open Learner Modelling. *International Journal of Artificial Intelligence in Education* 13(1), pp. 35-78.

S., Embretson, & S., Reise (2000). Item response theory for psychologists. Mahwah, NJ: Erlbaum.

A.C., Graesser, N., Person, D., Harter & TRG (2001). Teaching tactics and dialog in AutoTutor. *International Journal of Artificial Intelligence in Education*, 12, 257-279.

D., Hartley, & A., Mitrovic (2002). Supporting Learning by opening the Student Model. *In proceedings of ITS 2002.* pp. 453-462.

E., Horvitz, J.S., Breese, D., Heckerman, D., Hovel, & K., Rommelse (1998). The Lumiere Project: Bayesian user modeling for inferring the goals and needs of software users. *Fourteenth Conference on Uncertainty in Artificial Intelligence.* San Francisco: Morgan Kaufmann, 256-265.

W.L., Johnson, J.W., Rickel, & J.C., Lester (2000). Animated Pedagogical Agents: Face-to-Face Interaction in Interactive Learning Environments. *International Journal of Artificial Intelligence in Education*, 11(1), 47-78.

J., Kay (1998). A Scrutable User Modelling Shell for User-Adapted Interaction. Ph.D. Thesis, Basser Department of Computer Science, University of Sydney, Sydney, Australia.

F. M., Lord, & M. R., Novick, (1968). Statistical theories of mental test scores. Reading, MA: Addison-Wesley.

R. J., Mislevy, & D. H., Gitomer (1996). The Role of Probability-Based Inference in an Intelligent Tutoring System.*User Modeling and User-Adaptive Interaction. Special Issue on Numerical Uncertainty Management in User and Student Modeling*, vol 5(3), 253-282.

R., Nkambou, Y., Laporte, R., Yatchou, & G., Gouardres (2003). Embodied emotional agent in intelligent training system. *In Recent Advances in intelligent Paradigms and Applications*, J. Kacprzyk, A. Abraham, and L. C. Jain, (Eds.) Studies In Fuzziness And Soft Computing. Physica-Verlag Heidelberg, 235-253.

R.W., Picard (1997). Affective computing. Cambridge, MA: MIT Press.

J., Reye (2004). Student Modelling based on Belief Networks.*International Journal of Artificial Intelligence in Education.* Vol. 14, 63-96.

V.J., Shute, M., Ventura, M., Bauer & D., Zapata-Rivera (in press) Melding the Power of Serious Games and Embedded Assessment to Monitor and Foster Learning: Flow and Grow. To appear in L. Miller (Ed.), Serious Games: Learning, Development and Change.

K., VanLehn, & J., Martin (1997). Evaluation on an assessment system based on Bayesian student modeling. *International Journal of Artificial Intelligence in*

Education, Vol. 8. 179-221.

J.D., Zapata-Rivera, & J., Greer (2001). Visualizing and Inspecting Bayesian Belief Models. *Workshop on Effective Interactive AI Resources. International Joint Conference on Artificial Intelligence IJCAI 2001*, 47-49.

J.D., Zapata-Rivera, & J., Greer (2004). Interacting with Bayesian Student Models. *International Journal of Artificial Intelligence in Education.* Vol. 14, Nr. 2 127-163.

J.D., Zapata-Rivera, W., VanWinkle, Shute, V.J., J.S., Underwood & M.I. Bauer (2007). English ABLE. In Proceedings of the 13th International Conference of Artificial Intelligence in Education. 8pp.