Kerly, A., Hall, P. & Bull, S. (2006). Bringing Chatbots into Education: Towards Natural Language Negotiation of Open Learner Models, in R. Ellis, T. Allen & A. Tuson (eds), Applications and Innovations in Intelligent Systems XIV - Proceedings of AI-2006, 26th SGAI International Conference on Innovative Techniques and Applications of Artificial Intelligence, Springer.

Bringing Chatbots into Education: Towards Natural Language Negotiation of Open Learner Models

Alice Kerly[1], Phil Hall[2] and Susan Bull[1]

[1] Electronic, Electrical and Computer Engineering, University of Birmingham, Edgbaston, Birmingham, B15 2TT, UK
[2] Elzware Ltd, 70 Mendip Road, Windmill Hill, Bristol, BS3 4NY
[1] {alk584, s.bull}@bham.ac.uk [2] phil.hall@elzware.com

Abstract. There is an extensive body of work on Intelligent Tutoring Systems: computer environments for education, teaching and training that adapt to the needs of the individual learner. Work on personalisation and adaptivity has included research into allowing the student user to enhance the system's adaptivity by improving the accuracy of the underlying learner model. Open Learner Modelling, where the system's model of the user's knowledge is revealed to the user, has been proposed to support student reflection on their learning. Increased accuracy of the learner model can be obtained by the student and system jointly negotiating the learner model. We present the initial investigations into a system to allow people to negotiate the model of their understanding of a topic in natural language. This paper discusses the development and capabilities of both conversational agents (or chatbots) and Intelligent Tutoring Systems, in particular Open Learner Modelling. We describe a Wizard-of-Oz experiment to investigate the feasibility of using a chatbot to support negotiation, and conclude that a fusion of the two fields can lead to developing negotiation techniques for chatbots and the enhancement of the Open Learner Model. This technology, if successful, could have widespread application in schools, universities and other training scenarios.

1 Background

This paper unites work in chatbots, natural language processing in an educational context, intelligent tutoring systems and learner modelling. We briefly introduce these fields below, and propose how this combined approach might be used to support learners.

1.1 Chatbots

Conversational agents, or chatbots, provide a natural language interface to their users. Their design has become increasingly sophisticated and their use adopted in education, (e.g. [1]), commerce (e.g. [2], [3]), entertainment (e.g. [4]) and the public sector (e.g. [5], [6]).

ELIZA [7], was regarded as one of the first chatbots. ELIZA analysed input sentences and created its response based on reassembly rules associated with a decomposition of the input. This produced an impression of caring about its users, but it held no memory of the conversation and so could not enter into any form of targeted collaboration or negotiation. The syntactic language processing used by ELIZA has been developed significantly, leading to the development of a number of language processing chatbots (an exhaustive list can be found at [8]).

A.L.I.C.E. [9] is a chatbot built using Artificial Intelligence Markup Language (AIML), developed over the past 10 years. The chatbot is based on categories containing a stimulus, or pattern, and a template for the response. Category patterns are matched to find the most appropriate response to a user input. Further AIML tags provide for consideration of context, conditional branching and supervised learning to produce new responses. A.L.I.C.E. is a viable and experienced system but has not to our knowledge, as yet, been applied in a commercial environment.

The Jabberwacky [10] chatbot has as its aim to "simulate natural human chat in an interesting, entertaining and humorous manner". Jabberwacky learns from all its previous conversations with humans. It functions by storing everything that is said to it, and uses contextual pattern matching techniques to select the most appropriate response. It has no hard-coded rules, instead relying entirely on previous conversations. It is explicitly not intended to do anything 'useful', instead being simply to chat [10].

Modern commercial chatbots, such as those developed with Lingubot™ [11] technology, offer sophisticated development environments allowing the building of intelligent conversational agents with complex, goal driven behaviour. In 'Lingubots' both the words and the grammatical structure of the user's input are analysed using customised templates. This facilitates the development of a user model, which is used in conjunction with the conversational context and specific words in the dialogue to determine the chatbot's response. Responses might include further conversation with the user, reading or writing to external systems (for instance to open a web page or update a database), or a combination of these. This rich range of responses allows for intelligent conversation with the user, and provides the ability to steer the user back to the task in hand if they stray from the designated discussion content for too long.

As computing technology and the underlying language processing software progresses, we can expect to see potentially exponential growth in the delivered complexity of chatbots. Already, they have come a long way from their roots in systems that were more about fun, flirtation or simple 'chat'. We are now approaching a time where the technologies such as Lingubot can, through extensive syntactic structures developed for natural language processing and some complex methodological data structuring, begin to display behaviour that users will interpret as understanding.

1.2 Intelligent Tutoring Systems

The field of Intelligent Tutoring Systems emerged from earlier work on generative computer-assisted instruction, for example Uhr's [12] work on generating arithmetic problems. Other systems were able to adaptively select problems based on the student's performance (e.g. Suppes, 1967, cited by Sleeman and Brown [13], pg 1). These systems maintained basic models of the student's behaviour, but did not tend to store representations of the student's actual knowledge [13]. Uhr advocates systems that were able to generate new problems according to a small set of axioms, in order to provide problems that were suited to the level the learner was performing at [12]. Sleeman and Brown also argued that to tutor well the system must constrain the student's instructional paths by a system of student modelling [13].

Intelligent Tutoring Systems (ITS) researchers were able to exploit developments within both the cognitive sciences and in hardware to produce systems which took into account the learner's state, e.g. Clancey's GUIDON [14] and Burton's DEBUGGY [15] systems. There are a variety of learner modelling techniques, such as overlay models which model the learner as a subset of the expert; perturbation models that also allow misconceptions to be modelled; Bayesian networks to allow more complex inferences (see [16] for an overview).Work in learner modelling has continued to be central to research in intelligent tutoring systems (e.g. [17], [18], [19]), with researchers exploring issues such as learner control over the learner model contents, modelling learner misconceptions, peer-group modelling, presentation of models, and learner models for mobile computing.

Thus learner modelling has developed as the practice of creating a model of the learner's understanding based on their interaction with an ITS. This allows for personalisation of the user experience, and to provide individualised feedback to the user on their progress.

1.3 Learner Modelling and Open Learner Modelling

Intelligent Tutoring Systems employ a learner model to infer the learner's knowledge and to provide an adaptive interaction. While many ITSs do not reveal the contents of the learner model to the learner, it has been argued that opening the learner model to the ITS users can in fact provide opportunities for learner reflection and deep learning that enhances the learning experience (e.g. [20], [21], [22], [23] and [24]).

Open learner models are therefore accessible to the user. They are inferred from the learner's interaction with the system (as in any ITS), and may also include contributions obtained directly (explicitly) from the student. As a pedagogical goal, learner reflection is endorsed by many theories, including Dewey [25], Schön [26], and Kolb [27]. Bull & Pain [28] and Dimitrova [21] proposed that both learner reflection and model accuracy could be increased through a process of negotiation of the learner model contents and implemented the Mr. Collins and STyLE-OLM systems respectively. In this method the learner model is collaboratively constructed and maintained by both the system and the learner. In both the above systems, the learner was required to discuss their beliefs with the system, arguing against the system's assessment if they disagreed with it, and providing supporting evidence or argument for their own beliefs when they differed from the system. This interaction supported the increased learner reflection intended to benefit learning, and produced a more accurate learner model on which to base system adaptivity.

In order to support the negotiation functionality, the learner model must store distinct records of the learner's and the system's beliefs about the learner's knowledge. Two separate belief measures were maintained in the Mr. Collins [28] system, each of which was taken into account by the system in providing adaptive interactions. Baker's notion of interaction symmetry [29] was applied to the system, ensuring that all dialogue moves necessary for negotiation were available to both the student and the system. Laboratory studies of the Mr. Collins system [28] found that students were interested in being able to see the contents of their learner model. They were keen to use negotiation to improve the accuracy of the learner model and most students also wanted the system to challenge them if it disagreed with their attempts to change their confidence in their performance.

Previous open learner model systems have employed menu-selection or conceptual graphs to achieve negotiation of the learner model contents. While laboratory trials ([28], [21]) of these systems suggested the potential for engaging learner reflection and enhancing the accuracy of the learner model, the negotiation methods used may be restrictive or unnatural. We propose that natural language negotiation through a Chatbot may offer users the flexibility to express their views in a naturalistic and intuitive way. For the early exploratory design of the Chatbot structure we have opted to follow the negotiation strategies provided in Mr. Collins as these offer a structured and limited architecture on which to base our protocols.

1.4 Intelligent Tutoring Systems that use Natural Language

Intelligent tutoring systems that use natural language have largely tended to be either tutors, pedagogical agents or avatars. Avatars provide an engaging, personalised and simple interface, often (though not always) with an animated human character (e.g. an avatar to improve child users' engagement in a web-based game to teach home and city safety [30]).

Pedagogical agents are autonomous agents that occupy computer learning environments and facilitate learning by interacting with students or other agents (e.g. [31], [32], [33], and [34]). They may act as peers, co-learners, competitors, helpers or instructors. For effective pedagogy, agents should be able to ask and respond to questions, give hints and explanations, monitor students and provide feedback [35]. Tutoring may be provided by pedagogical agents, avatars, or other simpler mechanisms (e.g. staged textual hints) and is tailored to the individual learner to help them progress through the immediate task.

This paper presents a study involving a simulated chatbot to investigate combining the benefits of natural language interaction with negotiation of the learner model. This is a new direction for the use of natural language in ITS, drawing on work on pedagogical agents, avatars and natural language tutoring.

2 Using a Chatbot for Negotiated Learner Modelling

Kay [22] states that open learner models may enhance the student experience by encouraging effective learning (rather than merely browsing), by creating opportunities for the learner to disagree or negotiate with the system, by asking the student to reflect on their knowledge and compare this with their student model, and by asking the student to use their model to identify areas to revise. Kay also suggests that offering different presentations of the same information may help students to think about their knowledge in different ways. It is envisaged that all of these goals could be facilitated by a chatbot suitably integrated into an open learner model system.

Natural language dialogue has not previously been used in an ITS to support the creation or maintenance of an open learner model. Given the benefits that negotiation can bring to learner model construction and to encouraging learner reflection, and the capabilities of modern chatbots, the present study aimed to explore the feasibility of integrating a natural language conversational agent into an open learner model system to enable student-system negotiation over the contents of the learner model.

2.1 The Choice of Chatbot

The technology for negotiation of an open learner model requires the following characteristics:

- Keeping the user 'on topic' – preventing them from deviating too long from the issue at hand
- Database connectivity – to allow reading and writing of data to and from the learner model
- Capability to be event driven by database changes – to allow the chatbot to initiate negotiation if there are conflicts in the learner model
- Web integration – to allow easy deployment to maximum possible users
- An appropriate corpus of semantic reasoning knowledge.

The open source foundations of AIML [9] have provided an interesting and useful application for AI development. We believe that while the A.L.I.C.E. solution could be used to create this system, writing of AIML on this scale would be challenging for the non-developer, and that the richness of processing is limited in comparison to Lingubots [11]. Both Lingubots and AIML have a wide corpus, but Lingubot's is more focussed on goal driven conversation. The Lingubot Creator editorial interface allows for more in-depth developmental objectives and facilitates the building of this complex system.

The Lingubot technology has generated a significant corpus of external scripts and applications, providing functionality relevant to conversations, such as abuse counters, theme metrics and directed conversational abilities, as well as integration with many web technologies. It has the capability to generate and manipulate variables and information regarding the conversation, and also for retrieving and posting information to other web enabled applications such as search engines and databases. The commercial nature of the technology also ensures that it is well tested.

Lingubot technology is the most appropriate AI technology for this development due to the ability for the human to remain in the loop due to its ease of update and open reporting structures. It also meets the requirements of a user-centred experience, easy web integration and database connectivity. By allowing the AI researcher to focus more on the delivery or outcomes of conversation, rather than the underlying pattern matching effectiveness we hope to build a bridge towards a more complex understanding of the dynamics of humans with machines.

3 Wizard-of-Oz study

3.1 The Wizard-of-Oz Paradigm

In this study, as a precursor to development of a chatbot, we used the Wizard-of-Oz method (discussed below) to conduct an experiment where the role of the chatbot was taken by the experimenter; the 'Wizard'. The fact that the chatbot was actually a person was not revealed to the participants until after the study to ensure that data collected from their interaction would be pertinent to human-computer conversational design.

Wizard-of-Oz experiments are studies where participants are told that they are interacting with a computer system, when in fact the interaction is mediated by a human operator, the 'Wizard' [36]. The dialogues required to negotiate over a learner model are expected to be complex and varied in terms of language, tasks, and domain content [37]. The Wizard-of-Oz method has been shown to be appropriate for collecting data about user interaction in such complex domains [38] cited by [39]. The approach is also valuable in that computer-mediated human-human interaction data can be an unreliable source of information for some important aspects of human-computer dialogue design as humans and computers can be expected to perform differently in conversation [38]. As a technique it is good at eliciting application-specific linguistic characteristics, and is also significant as keyboard interaction appears to be a mode that alters the normal organization of discourse [36], thus rendering non-computer mediated data less valuable.

The Wizard-of-Oz paradigm necessarily involves deception of the participants about the nature of the experiment. This study was conducted within British Psychological Society Ethical Principles for Research with Human Participants guidelines [40], including debriefing participants by explaining the purposes of the deception and answering their questions or concerns about the process.

3.2 The Learner Modelling System

In order to investigate the potential for negotiation over the learner model, we needed an open learner model system capable of independently storing the system's and student's beliefs. The system employed was a modified version of Flexi-OLM [41], an open learner model in the domain of C programming. The system infers its beliefs from students' answers to multiple-choice and short-answer questions, and computes and displays the learner model as a result, in a choice of 7 formats to suit the learner's preferences for how to access the information: 2 knowledge map views; 2 hierarchical structures; 2 list views; and a textual description. Figure 1 shows the concept map view as an example, with the colour of the nodes portraying the knowledge level of each concept. The learners were also able to record their own beliefs about their knowledge, and this data was stored in parallel with the system's inferences, and could be viewed in the same 7 formats as the system's beliefs, for direct comparison. These two belief sets (system inferences and student-provided information about their beliefs) provided the necessary data for comparisons of beliefs leading to negotiation of the learner model.

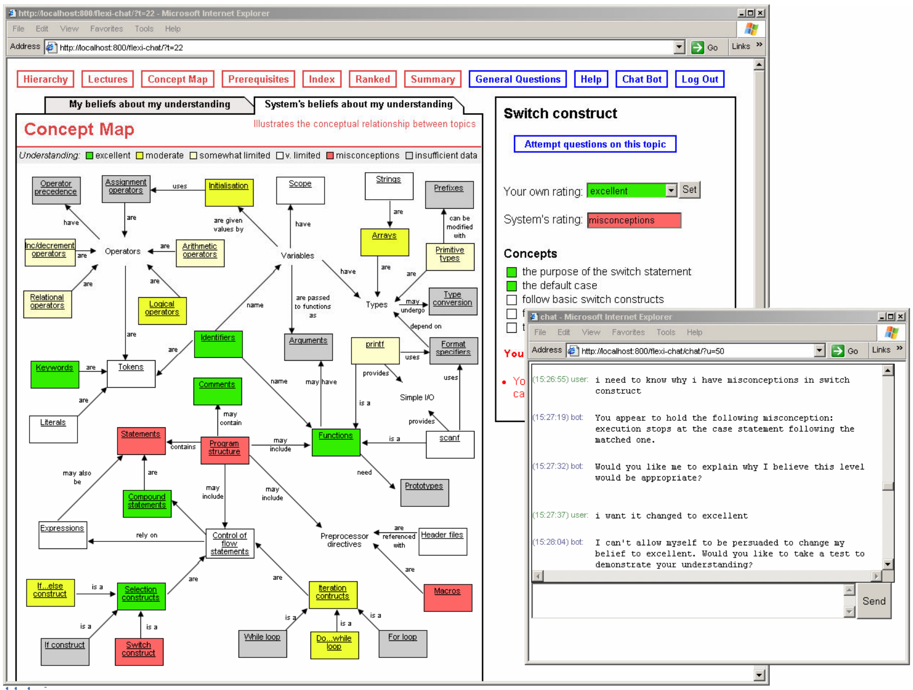

Figure 1. User's screen in Concept Map view of System beliefs of student knowledge & chat interface [41]

3.3 Experimental setup

The participants were 30 students from the University of Birmingham Electronic, Electrical and Computer Engineering Department. 11 were final year undergraduates and 19 were MSc students. All had previously taken courses in educational technology and C programming. All were competent English language speakers, though in some cases English was not their first language.

Mr. Collins [28] provided a range of strategies for negotiation: ask user if they wish to accept the system's viewpoint; offer compromise; ask user to justify their belief (e.g. by taking a test to demonstrate knowledge); system justify its belief; or offer student the opportunity to view the learner model. These strategies were adopted as the initial conversational basis of the ‘chatbot’. The ‘Wizard’ was provided with a decision tree to allow the consistent selection of appropriate responses, and 350 pre-authored ‘chatbot’ negotiation initiations and responses to user inputs [37]. These ‘canned responses’ can be seen in the left part of Figure 2, while the right shows the wizard’s view of the learner model. This view of the model enabled the wizard: (i) to compare the user and system beliefs for each topic in two columns using colour to represent the student's and system's beliefs about the learner's knowledge level; (ii) to see the student’s answers to questions; and (iii) to select unanswered questions to offer the student a test by which to demonstrate their knowledge.

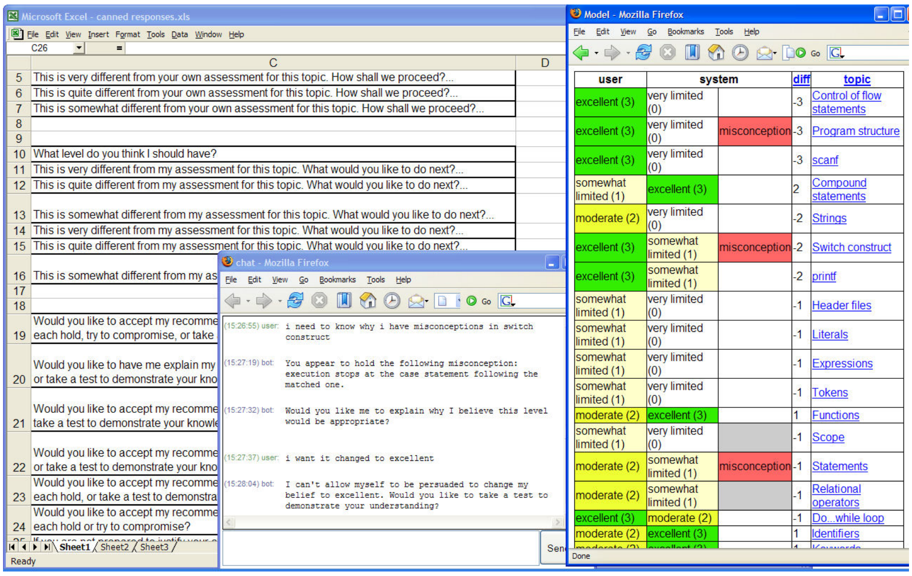

Figure 2. Wizard's view showing (left) part of the canned responses; (centre) Chat interface - excerpt reproduced below; (right) part of the Learner Model

- Chatbot: Would you like me to explain why I believe this level would be appropriate?
- User: I want it changed to excellent.
- Chatbot: I can’t allow myself to be persuaded to change my belief to excellent. Would you like to take a test to demonstrate your understanding?

The users were instructed that they would be able to use a chatbot feature of Flexi-OLM to discuss their learning of C programming which could lead to modification of the learner model. They were not constrained as to how they should interact with the ‘chatbot’. They were provided with an interface to allow them to answer questions related to C programming; view the system’s model of their knowledge or their own belief model; attempt to edit their own beliefs (though the ‘chatbot’ may challenge them) and negotiate over the model by challenging the chatbot about the system’s beliefs. Students completed post-interaction questionnaires with responses required on a 5-point Likert Scale. All interactions were logged by the system.

It is acknowledged that providing the Wizard with 350 pre-authored phrases substantially restricted the content that the 'chatbot' could handle. However, this limitation was necessary in order to remain manageable for the Wizard. We also recognise the limitations of the relatively small population, their restricted demographic, and their prior experience of educational technology.

3.4 Outcomes

In the 30 logged student-'chatbot' conversations there were 164 negotiation fragments (elements of conversation where a topic is introduced, discussed and resolved). Users initiated the conversation on 94 occasions; the 'chatbot' began the discussion 70 times. Interaction times ranged from 18 to 64 minutes, (mean 30.17 minutes, median 29.5 minutes). Negotiation fragments lasted a mean of 4.65 minutes (median 4 minutes).

Table 1 User opinions of the Chatbot and negotiation facility

	< strongly agree....strongly disagree >					Mean score
	(1)	(2)	(3)	(4)	(5)	
The negotiation changed my view of my understanding	6	13	8	3	0	2.3
I used the Chatbot to help me understand my learner model *	7	16	4	1	1	2.1
I always challenged the Chatbot when I disagreed with it (or I would)	13	11	5	1	0	1.8
I liked the Chatbot when it disagreed with me	4	16	6	2	2	2.4
I liked the Chatbot when it agreed with me *	6	16	7	0	0	2.0
I enjoyed interacting with the Chatbot *	12	11	6	0	0	1.8
I was able to achieve the negotiation tasks I wished with the Chatbot	6	18	5	0	1	2.1
The Chatbot made negotiation easy	7	17	4	1	1	2.1
The Chatbot was a convenient way to provide my opinions to the system *	8	14	6	1	0	2.0
It was easy to interact with the Chatbot	7	13	7	2	1	2.2

(* indicates one user did not respond, therefore total responses is 29)

Table 1 summarises the questionnaire responses. For each of the above statements the mean score shows that users were either positive or neutral in their reaction. Particularly noteworthy are the 13 strongly agree and 11 agree responses to "I always challenged the Chatbot when I disagreed with it" as this willingness to engage in discussion is essential to the concept of negotiating a learner model. Also of interest are the 12 strongly agree and 11 agree responses to "I enjoyed interacting with the Chatbot". This enjoyment in the interaction is crucial in keeping users engaged and involved. Interestingly, most students claimed to like the chatbot even when it disagreed with them. It is also important that 17 users agreed and 7 strongly agreed that "the Chatbot made negotiation easy". It is our aim to provide a user-friendly interface to negotiation, and this finding supports our proposal that chatbots are a fitting tool in this endeavour. Of key interest to the pedagogical aims of this work are the first two statements. 13 students agreed and 6 strongly agreed that "the negotiation changed my view of my understanding" and 16 agreed and 7 strongly agreed that "I used the Chatbot to help me understand my learner model". There were also positive results for statements regarding usability of the chatbot, including being able to achieve the negotiation tasks wished, being a convenient way to provide opinions to the system, and ease of interaction.

Our investigations suggested that allowing students to additionally discuss content not related to their learner model may be useful in helping them to build a rapport with the chatbot. This is important in keeping students engaged in the discussion for as long as possible and to permit the maximum possible on-topic negotiation. Students in the Wizard-of-Oz study commented that they would like to be able to "discuss other stuff, e.g. weather, just for a little break from working!" Users also often asked the chatbot "how old are you?" and "what is your name?" One user asked "what will you do this evening?" and several others asked "what's the time?"

From the transcripts of the Wizard-of-Oz study we were able to identify a number of questions that frequently arose in users' conversations with the chatbot. Therefore, we believe that in a learner modelling environment, key inputs that the chatbot should be able to respond to include:

- What should I do next?
- Can you help me?
- What should I study?
- What am I good at?
- What am I bad at?
- What's the answer?
- What are my misconceptions?

These findings suggest that implementation of a chatbot to facilitate negotiation of the learner model is a worthwhile next step in open learner modelling.

4 Lessons for implementation

As a result of our investigations, we have established a set of requirements for a system that is to provide negotiation functionality, particularly within a learner modelling environment.

4.1 System Requirements

- **Links to external databases**. This is essential in order to provide accurate and current modelling capabilities. The chatbot must be able to write to the database, for example to record the outcomes of a negotiation fragment, and must also be able to read data from the database in order to insert into its text outputs, for example to ask the user a test question.

- **Common user requests.** The chatbot should be able to respond appropriately to requests identified as frequent or of particular importance.
- **Privacy of data**. The focus here is on reassuring users that data is secure.
- **Keeping the user on topic**. This is a particular skill of the Lingubot technology, and is important for this system to be pedagogically successful. While a degree of smalltalk is considered beneficial to the interaction (see below) the essence of the system is negotiation over the learner model. If the learner is allowed to be distracted from their educational goals, this will be detrimental to the success of the learning episode. Therefore, the chatbot must be able to manage a small amount of "extra-curricular" discussion before pleasantly but firmly returning the user to the topic at hand.
- **The need for smalltalk.** We intend to use the power of the Yhaken conversational core (a specific Lingubot implementation) [42] to allow the user an effective smalltalk experience and deal with flak that arises when we push the user back to the point at hand.
- **Prevent user 'losing' chatbot.** In order to ensure the user always has access to the chatbot's conversation, it is necessary to ensure that they cannot have accidentally closed or buried the chatbot window. We propose overcoming this risk by embedding the chatbot window in a frame at the side of the learner modelling system. This will ensure it remains accessible at all times, and that users may better associate the chatbot's utterances with their behaviour and results in the learner model.
- **Understand negotiation fragments.** There is a need for the chatbot to 'know' when a negotiation fragment has been completed. This is necessary to ensure that discussions remain on target, and that the chatbot can address further outstanding issues once one has been dealt with satisfactorily.
- **Deliver an effective conversation.** We believe that this will be achieved by ensuring that the user is presented with self-referential objectives (i.e. always referring to pertinent topics from their learner model) and pushed to the edges of understanding. As a piece of educational software, this type of behaviour aims to promote the reflection that has been shown [33] to affect deep learning.
- **Feedback mechanism.** To ensure continual improvement of the system it will be necessary to incorporate methods to review the transcripts of conversations, assess their success, and make additions or alterations to the Lingubot scripts or functions as necessary.

4.2 Areas for Further Research

We have also identified a number of areas which will benefit from further research, but which are (at least initially) non-essential to the system.

- **What evidence the system can accept.** Currently the only way for users to convince the chatbot and system of their understanding of a subject is to correctly answer test questions. It would be interesting to explore what other evidence may be acceptable, for example details of courses, work, training or experience through which the user is able to demonstrate the relevant knowledge point.

- **Graphical versus text display.** In our initial implementation, the chatbot is likely to be text only, perhaps represented by a static graphic. Many chatbots and pedagogical agents use “embodied” or graphical displays. An investigation of the effect on user perceptions of the system when the chatbot is text only or has a face/character/graphical display of “emotions”, designed according to findings from other educational environments, will further benefit our understanding of user perceptions and rapport in the context of discussing one's knowledge in negotiated learner modelling.

5 Conclusion

We have presented an investigation into the potential for two new developments: the use of chatbots to provide negotiation facilities, and the incorporation of chatbots into an open learner modelling environment.

Using a Wizard-of-Oz experiment, we investigated the requirements for constructing a chatbot in a negotiated open learner model, the potential utility of a chatbot in such a scenario, and user reactions to the interaction. We established technical constraints, including database access, data privacy and a feedback mechanism, and requirements for effective dialogue, including keeping the user on-topic, incorporating smalltalk capabilities, preventing crossed dialogue threads, and successfully resolving negotiation fragments. Users in the study commented on their enjoyment of the interaction, the ease of negotiation by this method, and their willingness to engage with a chatbot, thereby improving their learner model.

We conclude that, while there are technical challenges to be met in specifying the scripts and functions of a complete system, a chatbot can provide the necessary negotiation facilities for an enhanced open learner model interaction. We have commenced work on the detailed specification and build of the system, and intend to publish further results on user interaction with our ‘live’ chatbot.

Acknowledgements

We thank Andrew Mabbott for extending Flexi-OLM [41] for this study, and Felix Cohen for his comments on draft versions of this paper. The first author is funded by the UK Engineering and Physical Sciences Research Council.

References

1. Jia, J., CSIEC (Computer Simulator in Educational Communication): An Intelligent Web-Based Teaching System for Foreign Language Learning, in *ED-MEDIA (World Conference on Educational Multimedia, Hypermedia & Telecommunications)*. 2004: Lugano, Switzerland.
2. De Angeli, A., G.I. Johnson, and L. Coventry, The unfriendly user: exploring social reactions to chatterbots, in *International Conference on Affective Human Factors Design*, Helander, Khalid, and Tham, Editors. 2001, Asean Academic Press: London.
3. Creative Virtual. UK Lingubot Customers. 2004-2006. Listing of major companies using Linubot technology. www.creativevirtual.com/customers.php. Accessed 14/12/05

4. Wacky Web Fun Ltd. RacingFrogz.org. c2005. Online game with chatbot capabilities. www.racingfrogz.org. Accessed 17/05/06
5. West Ham and Plainstow NDC. 2005. New Deal For Communities. "Splodge" New Deal For Communities chatbot assistant. www.ndfc.co.uk. Accessed 02/06/06
6. Bath & North East Somerset Council. 2006. Council web site Citizen Information Service Agent. www.bathnes.gov.uk/BathNES/Search/ask/default.htm. Accessed 02/06/06
7. Weizenbaum, J., ELIZA - A Computer Program For The Study of Natural Language Communications Between Man and Machine. *Communications of the ACM*, 1966. **9**(1): p. 36-45.
8. dmoz. 2006. Open Directory Project Computers: Artificial Intelligence: Natural Language: Chatterbots http://dmoz.org/Computers/Artificial_Intelligence/Natural_Language/Chatterbots/ Accessed 08/06/06
9. Wallace, R.S. Chapter 00. The Anatomy of A.L.I.C.E., in http://www.alicebot.org/documentation. Accessed 09/12/05
10. Carpenter, R. 1997-2006. Jabberwacky.com. www.jabberwacky.com. Accessed 16/05/06
11. Creative Virtual. 2004-2006. CreativeVirtual.com. Web site of UK Lingubot distributor. www.creativevirtual.com. Accessed 12/05/06
12. Uhr, L., Teaching machine programs that generate problems as a function of interaction with students, in *24th National Conference*. 1969. p. 125-134.
13. Sleeman, D.H. and J.S. Brown, eds. *Intelligent Tutoring Systems*. 1982, Academic Press: London.
14. Clancey, W.J., Tutoring Rules for Guiding a Case Method Dialogue, in *Intelligent Tutoring Systems*, D.H. Sleeman and J.S. Brown, Editors. 1981, Academic Press: London. p. 201-225.
15. Burton, R.R., Diagnosing Bugs in Simple Procedural Skills, in *Intelligent Tutoring Systems*, D.H. Sleeman and J.S. Brown, Editors. 1981, Academic Press: London. p. 157-183.
16. Holt, P., S. Dubs, M. Jones, and J. E. Greer, The State of Student Modelling, in *Student Models: The Key to Individualized Educational Systems*, J.E. Greer and G.I. McCalla, Editors. 1994, Springer-Verlag: New York.
17. Michaud, L.N. and K.F. McCoy, Capturing the Evolution of Grammatical Knowledge in a CALL System for Deaf Learners of English. *International Journal of Artificial Intelligence in Education,* 2006. **16**(1): p. 65-97.
18. Zapata-Rivera, J.D. and J.E. Greer, Interacting with Inspectable Bayesian Student Models. *International Journal of Artificial Intelligence in Education,* 2004. **14**(2): p. 127-163.
19. Reye, J., Student modelling based on belief networks. *International Journal of Artificial Intelligence in Education,* 2004. **14**: p. 63-96.
20. Bull, S., P. Brna, and H. Pain, Mr. Collins: A collaboratively constructed, inspectable student model for intelligent computer assisted language learning. *Instructional Science*, 1995. **23**: p. 65-87.
21. Dimitrova, V., STyLE-OLM: Interactive Open Learner Modelling. *International Journal of Artificial Intelligence in Education*, 2003. **13**: p. 35-78.
22. Kay, J., Learner Know Thyself: Student Models to give Learner Control and Responsibility, in *ICCE97 International Conference on Computers in Education*, Z. Halim, T. Ottomann, and Z. Razak, Editors. 1997: Kuching, Malaysia. p. 17-24.
23. Morales, R., Exploring participative learner modelling and its effects on learner behaviour. 2000, Unpublished PhD Thesis. University of Edinburgh: Edinburgh.
24. Zapata-Rivera, J.D. and J. Greer, Externalising Learner Modelling Representations, in *Workshop on External Representations of AIED: Multiple Forms and Multiple Roles.*

International Conference on Artificial Intelligence in Education. 2001: San Antonio, Texas.
25. Dewey, J., *How We Think: A Restatement of the Relation of Reflective Thinking to the Educative Process*. 1933, Boston: D. C. Heath and Company.
26. Schön, D., *The Reflective Practitioner*. 1983, London: Maurice Temple Smith Ltd.
27. Kolb, D., The Process of Experiential Learning, in *The Experiential Learning: Experience as the Source of Learning and Development*, D. Kolb, Editor. 1984, Prentice-Hall: NJ.
28. Bull, S. and H. Pain, Did I say what I think I said, and do you agree with me?: Inspecting and Questioning the Student Model, in *World Conference on Artificial Intelligence in Education*, J. Greer, Editor. 1995, AACE: Charlottesville, VA. p. 501-508.
29. Baker, M.J., Negotiated Tutoring, an Approach to Interaction in Intelligent Tutoring Systems. 1990, Unpublished PhD Thesis. Open University: Milton Keynes, UK.
30. Sheth, R., Avatar Technology: Giving a Face to the e-Learning Interface. *E-Learning Developer's Journal*, 2003 (August 25 2003).
31. André, E., T. Rist, and J. Müller. WebPersona: A Life-Like Presentation Agent for Educational Applications on the World-Wide Web. in *Workshop on Intelligent Educational Systems on the World Wide Web at Artificial Intelligence in Education*. 1997. Kobe, Japan.
32. Shaw, E., W.L. Johnson, and R. Ganeshan, Pedagogical Agents on the Web, in *International Conference on Autonomous Agents*. 1999, ACM Press: Seattle, WA, USA. p. 283-290.
33. Graesser, A.C., K. Moreno, J. Martineau, A. Adcock, A. Olney, and N. Person. AutoTutor improves deep learning of computer literacy: Is it the dialogue or the talking head? in *Artificial Intelligence in Education*. 2003. Amsterdam: IOS Press.
34. Lester, J.C., B.A. Stone, and G.D. Stelling, Lifelike Pedagogical Agents for Mixed-Initiative Problem Solving in Constructivist Learning Environments. *User Modeling and User-Adapted Interaction,* 1999. **9**: p. 1-44.
35. Person, N.K. and A.C. Graesser, Pedagogical Agents and Tutors, in *Encyclopedia of Education*, J.W. Guthrie, Editor. 2006, Macmillan: New York. p. 1169-1172.
36. Dählback, N., A. Jönsson, and L. Ahrenberg. Wizard of Oz Studies - Why and How. in *International Workshop on Intelligent User Interfaces '93*. 1993: ACM.
37. Kerly, A. and S. Bull, The Potential for Chatbots in Negotiated Learner Modelling, in *Intelligent Tutoring Systems: 8th International Conference*, M. Ikeda, K. Ashley, and T.-W. Chan, Editors. 2006, Springer Verlag: Berlin-Heidelberg. p. 443-452.
38. Bernsen, N.O., H. Dybkjaer, and L. Dybkjaer, *Designing Interactive Speech Systems – From First Ideas to User Testing.* 1998: Springer. New York.
39. Fiedler, A. and M. Gabsdil, Supporting Progressive Refinement of Wizard-of-Oz Experiments, in *Workshop on Empirical Methods for Tutorial Dialogue Systems at International Conference on Intelligent Tutoring Systems*, C. Penstein Rose and V. Aleven, Editors. 2002: San Sebastian, Spain. p. 62-69.
40. British Psychological Society. 2000-2004.BPS Ethical Principles for Research with Human Participants. http://www.bps.org.uk/the-society/ethics-rules-charter-code-of-conduct/code-of-conduct/ethical-principles-for-conducting-research-with-human-participants.cfm Accessed 18/10/05
41. Mabbott, A. and S. Bull, Student Preferences for Editing, Persuading and Negotiating the Open Learner Model, in *Intelligent Tutoring Systems: 8th International Conference*, M. Ikeda, K. Ashley, and T.-W. Chan, Editors. 2006, Springer Verlag: Berlin-Heidelberg. p. 481-490.
42. Elzware Ltd. Yhaken.com. Website for Yhaken, a Lingubot built with Kiwilogic technology. www.yhaken.com. Accessed 01/06/06