Conversational Agents in E-Learning

Alice Kerly[1], Richard Ellis[2], Susan Bull[3]

Abstract This paper discusses the use of natural language or 'conversational' agents in e-learning environments. We describe and contrast the various applications of conversational agent technology represented in the e-learning literature, including tutors, learning companions, language practice and systems to encourage reflection. We offer two more detailed examples of conversational agents, one which provides learning support, and the other support for self-assessment. Issues and challenges for developers of conversational agent systems for e-learning are identified and discussed.

1 Introduction

This paper provides an overview of the use of natural language or 'conversational' agents in educational systems. With the growing maturity of conversational technologies, the possibilities for integrating conversation and discourse in e-learning are receiving greater attention in both research and commercial settings. Conversational agents have been produced to meet a wide range of applications, including tutoring (e.g. [1],[2],[3]), question-answering (e.g. [4],[5],[6]), conversation practice for language learners, (e.g. [7],[8]), pedagogical agents and learning companions (e.g. [9],[10],[11],[12]), and dialogues to promote reflection and metacognitive skills (e.g. [13],[14]). Conversational agents build on traditional education systems, providing a natural and practical interface for the learner. They are capable of offering bespoke support for each individual, and to recognise and build upon the strengths, interests and abilities of individuals in order to foster engaged and independent learners.

[1] Alice Kerly
Electronic, Electrical and Computer Engineering, University of Birmingham, Edgbaston, Birmingham, B15 2TT, UK. alk584@bham.ac.uk

[2] Richard Ellis
Elzware Ltd, 70 Mendip Road, Bristol, BS3 4NY, UK. richard.ellis@elzware.com

[3] Susan Bull
Electronic, Electrical and Computer Engineering, University of Birmingham, Edgbaston, Birmingham, B15 2TT, UK. s.bull@bham.ac.uk

The paper has four further parts. Section two introduces conversational agents and their use in e-learning systems, describing and contrasting their various applications. Using these examples, we then consider the evidence for the value of using conversational agents in educational technology in section three. Section four introduces two extended examples of conversational agents in e-learning, and discusses their aims and development, before we identify and discuss issues for consideration by developers of conversational agents in educational systems in section five.

2 Conversational Agents in E-Learning

This section examines the literature from the artificial intelligence in education, intelligent tutoring systems and associated educational technology and e-learning fields, revealing a substantial research and development effort focussed on the application of conversational agents in educational software. We describe some of the key examples below and consider the strengths of natural language supported systems in section 3.

The design, implementation and strategies of conversational systems employed in e-learning varies widely, reflecting the diverse nature of the evolving conversational agent technologies. The conversations are generally mediated through simple text based forms (e.g. [2]), with users typing responses and questions at a keyboard. Some systems use embodied conversational agents (e.g. [1]) capable of displaying emotion and gesture, whereas others will use a simpler avatar (e.g. [14]). Voice output, using text to speech synthesis is used in some systems (e.g. [1]), and speech input systems are increasingly viable (e.g. [15]).

One of the most substantially investigated areas for the use of natural language dialogue in e-learning has been in providing tutoring. The AutoTutor project [1] is a notable example of this, providing tutorial dialogues on subjects including university level computer literacy and physics. The tutoring tactics employed by AutoTutor assist students in actively constructing knowledge, and are based on extensive analysis of naturalistic tutoring sessions by human tutors. The technology behind the system includes the use of a dialogue manager, curriculum scripts and latent semantic analysis. This system has been demonstrated to give a grade improvement of .5 to .6 SD units when compared to control conditions for gains in learning and memory.

Another tutoring system employing dialogue is Ms Lindquist [2], which offers ‘coached practice’ to high school students in algebra by scaffolding ‘learning by doing’ rather than offering explicit instruction. Early work with the system found that students using Ms Lindquist did fewer problems, but that they learned equally well or better than students who were simply told the answer. The results also suggested that the dialogue was beneficial in maintaining student motivation. The authors concluded that Ms Lindquist was a ‘Less is More’ approach, where learn-

ers tackled fewer problems, but learnt more per problem when they were engaged in an intelligent dialogue [2].

A final example of natural language tutoring is the Geometry Explanation Tutor [3] where students explain their answers to geometry problems in their own words. The system uses a knowledge-based approach to recognize explanations as correct or partially correct, and a statistical text classifier when the knowledge-based method fails. Studies with the Geometry Explanation Tutor [16] found that students who explain in a dialogue learn better to provide general explanations for problem-solving steps (in terms of geometry theorems and definitions) than those who explain by means of a menu.

The second major grouping we identify is that of pedagogical agents and learning companions. Learning companions are simulated characters that act as a companion to the student, and take a non-authoritative role in a social learning environment [17]. In general, a pedagogical agent will be a more authoritative teacher [17]. These agents, which may employ gesture, synthesised speech and emotional facial displays, have been investigated in domains ranging from helping children to learn plant biology [9] through to continuing medical education [10] and naval training [18]. Research into the roles which may be played by a pedagogical agent or learning companion has investigated agents as mentors, peers, experts or instructors [11]. In some systems the student must teach the agent [19], or interact with agent peers or co-learners [12], who may even include trouble makers intended to provoke cognitive dissonance to prompt learning [20]. Researchers have also investigated user preferences for level of agent expertise. Findings suggest that in general, similarities in competency between an agent and learner have positive impacts on learners' affective attainments, such that academically strong students showed higher self-efficacy beliefs in a task after working with a high-competency agent, while academically weak students showed higher self-efficacy after working with a low-competency agent [21].

Other research has examined the provision of conversational agents for question and answer. For example, TutorBot [4] simulates a didactic tutor and allows the user to retrieve information from a knowledge source using natural language in a question/answer schema. Other question-answering systems have included a student discussion board [5] where the conversational agent mines a corpus to retrieve an answer based on cosine similarities between the query post and the corpus passages, and the Intelligent Verilog Compiler Project [6] which allows learners to ask questions in English that query the same ontology as is used to provide the system's 'help' texts. This style of use most closely mirrors the most common use in commercial environments where conversational agents are used for information retrieval (see [22] for examples).

Conversational agent systems have been proposed to offer conversation practice for language learners. Jia [23] found that users were dissatisfied with the responses provided by a basic ALICEbot (www.alicebot.org) implementation, and the pattern-matching mechanism was deemed insufficient for use as a foreign language practice environment. In contrast, Jabberwacky [7]

(www.jabberwacky.com) uses a very different technology to ALICEbots, learning from all its previous conversations. It has been suggested for providing language practice; Fryer and Carpenter note that agents are willing to repeat the same material as often as students require. They also argue that chatbots give students the opportunity to use varying language structures and vocabulary (for example slang and taboo words) which they may otherwise get little chance to experience [7]. Chatbots trained on a corpus have been proposed to allow conversation practice on specific domains [8]. This may be restrictive as the system can then only 'talk' on the domain of the training corpus, but the method may be useful as a tool for languages which are unknown to developers or where there is a shortage of existing tools in the corpus language [8]. The chatbots may also be augmented with non-corpus content.

Finally, we note the use of dialogue as a prompt for reflection. Grigoriadou et al. [13] describe a system where the learner reads a text about a historical event before stating their position about the significance of an issue and their justification of this opinion. Answers are classified as scientific, towards-scientific or non-scientific, and a dialogue generator produces "appropriate reflective diagnostic and learning dialogue for the learner". CALMsystem [14] promotes reflection of a different kind. Users answer questions on the domain, and state their confidence in their ability to answer correctly. The system infers a knowledge level for the student based on their answers, and encourages the learner to engage in a dialogue to reflect on their self-assessment and any differences between their belief and that of the system about their knowledge levels. Studies have shown this dialogue improved self-assessment accuracy significantly more than reflection based only on visual inspection of the system and learner beliefs [24].

This section has demonstrated the wide variety in conversational systems in e-learning. Implementations may employ full embodied conversational agents with emotion or gesture display, synthetic voice output, simple text-based output, dialogue with an accompanying avatar, and many variants or combinations of these. Developers have integrated conversational capabilities into systems for a range of reasons. We now discuss some of these motivations and benefits.

3 Motivations for Using Natural Language in E-Learning

As shown in section 2, conversational agents may offer improvements in grades or other measurable learning gains. AutoTutor [1] was able to demonstrate improvement better than an (untrained) human tutor, while Ms Lindquist [2] showed efficiency benefits as users learnt more from fewer exercises. Motivation and user engagement enhancements have also been frequently noted (e.g. [2], [11], [7]). In some cases motivation may be actively supported through deliberate motivational tutoring techniques; in others it may be a useful by-product of exposure to a novel technique. If motivational benefits are to be retained, then this novelty cannot be

relied upon, and further research into deliberate scaffolding of affect may be required.

A key feature of conversational agents in e-learning is the use of a natural communication method; dialogue is a medium through which nearly all learners are familiar with expressing themselves. Proponents of using natural language in educational (and other) interfaces argue that the use of normal communicative behaviour allows users' cognitive resources to be spent on the learning task, rather than stretched by application in the communication medium [25]. Computer literacy, and familiarity with online chatting media, is becoming ubiquitous and a greater number of users are expected to find conversing with their e-learning tool a feasible option.

Systems may include an escalation mechanism, allowing conversations meeting certain criteria (for example, causes for concern) to be forwarded by email or SMS to teachers, either in real time, or in later summaries (e.g. in TeachBot, section 4). This allows teachers to offer further support where weaknesses are identified, and illustrates the power of using conversational systems in parallel with existing teaching methods.

Social learning enables students to interact with other people, such as teachers or fellow students (real or computer-simulated), during learning activities [17]. There may be possibilities to integrate conversational agent systems in Learning 2.0 communities, as assistants, moderators, guides or as virtual peers within the community. Dialogue and anthropomorphic characteristics of pedagogical and conversational agents may help support the social dimension of e-learning activities, and the social context has been argued to catalyze the cultivation of, and motivation for, knowledge [17].

Conversational agents may also support learning reinforcement in a particular activity – rather than or in addition to teaching new concepts. TeachBot (see section 4) helps students while undertaking classroom tasks or homework, with users carrying out elements of work (e.g. planning, note-making, drafting) alone, and then returning to the bot at intervals to receive further support. In combination with the provision of PDAs to individual learners, this approach allows learners to have a personalised supporter constantly beside them at their desk.

Natural language e-learning further offers learning reinforcement where it is used to encourage users to review self-assessments of their ability to tackle questions on a given topic [14]. It can help students review their actual knowledge, explaining and defending their understanding. These prompts to students to consider their understanding allow learning in previous contexts to be reinforced by the conversational interaction. Formative or summative assessments may also be derived from conversational agent dialogues, and may be used to deliver feedback to either students or teachers.

4 Examples

This section details two examples, selected to illustrate different uses of conversational agents in e-learning and blended learning. Both are built using the Lingubot [22] technology, but the aims and designs of the systems are very different. The first example (CALMsystem) is for a specific (and novel) use: the conversational agent offers students the opportunity to engage in negotiation, as an adjunct to a wider open learner model system. In contrast to this, the conversational agent in the second example (TeachBot) provides the principal interface for a range of learning support functions. Rather than focussing on tutoring or providing a specific learning task in itself, it helps students with their existing class or homework tasks.

Our first example is **CALMsystem** - the Conversational Agent for Learner Modelling. CALMsystem provides a browser-based environment which allows learners to view the learner model (the representation of their understanding of various topics traditionally used by an intelligent tutoring system to adapt to the learner's needs), created as they answer questions within the system. Whenever the learner answers questions they are required to state their confidence in their ability to answer questions on that topic, thereby creating a model of their own self-assessments for the topics which is held in parallel with the system's inferences about their knowledge. The opportunity to view their learner model in this way is intended to encourage self-directed learning, and the development of metacognitive skills.

The conversational agent offers learners an opportunity to discuss their learner model through a natural language interface. They are encouraged to discuss any differences between their own and the system's assessments of their knowledge, and to negotiate changes to the model through a range of strategies including accepting the system's recommended knowledge level, asking the system to explain or justify its beliefs, asking for further information from the model, attempting to reach a compromise with the system about their abilities, or answering further test questions in order to justify their own belief and provide additional evidence to the system. Negotiated learner modelling systems have previously been developed which utilised menu-selection [26] or conceptual graphs [27] to facilitate this negotiation. Laboratory trials of these systems suggested potential for engaging learner reflection, but the negotiation methods may have been restrictive or unnatural to users. CALMsystem proposes that offering negotiation (a new direction for conversational agents) via natural language conversation will be more flexible, naturalistic and intuitive for learners. Investigations have also indicated the benefit of allowing learners to have limited conversations which are not on-topic, i.e. 'smalltalk'. This appears to be valuable in helping users to engage and develop rapport with the agent [28].

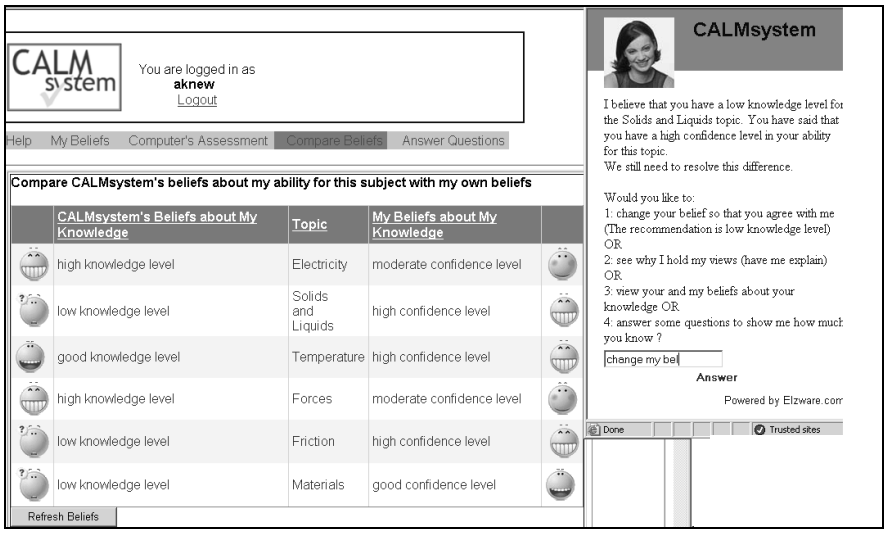

Figure 1 CALMsystem interface showing conversational agent in parallel with a browser display

The browser-based environment of CALMsystem operates independently of any intelligent tutoring system, thereby allowing access to users from a variety of platforms. It does not intend to be a teaching system, instead focussing on the benefits that arise from encouraging users to develop their metacognitive skills, improve their self assessment, and to develop as independent learners who can plan where they need to target their study. Studies with 55 UK Primary school children (aged 8-11) have found that children who interacted with CALMsystem including the conversational agent improved the accuracy of their self-assessments, and reduced the discrepancy between their own and the system assessments significantly more than children who used the environment without support from the conversational agent [14],[24].

The second example is **TeachBot**, a family of intelligent personalised learner support systems being developed by Elzware Ltd, in conjunction with Bristol (UK) City Learning Centres. TeachBots are designed for use by individual students to consolidate their knowledge of a subject and to act as a helper, providing advice in approaching and completing tasks such as writing an essay or designing a piece of furniture. In this way, TeachBot combines functions usually seen discretely in other conversational agents, such as tutored support by breaking questions into sub-questions (e.g. [2]), encouraging students to explain concepts in their own words (e.g. [3]), offering question-answering (e.g. [6]) and the social support of a learning companion (e.g. [17]).

TeachBot agents can be customised to provide support across a range of subjects, with previous versions supporting A-level (aged 16 - 18) Design and Technology students. The more advanced TeachBot implementation currently undergoing trials in Bristol schools has been developed to support Year 9 and 10

students (aged 13-15) studying for their GCSE in English Language. In addition to running on standard PCs, TeachBot has been developed for use on 3G connected handheld devices, providing readily accessible support to students both within and beyond the classroom.

The principal TeachBot interface, for user input and delivering material, is provided by the conversational agent, based, like the CALMsystem chatbot, on Lingubot technology. This is supplemented by a reactive avatar, showing emotions appropriate to the current conversation, and by access to a list of Frequently Asked Questions (FAQs) (see Figure 2).

Once logged in, the system guides the user through the appropriate aspects of their task, offering advice and challenging the user to consider different options and approaches. The notable difference between this and tutoring systems (e.g. [1],[2],[3]) that provide the learning objective and task or exercise is that TeachBot supports the user in their current activity from outside the system. For example, it may guide a student through critiquing the choice of materials in a design project, or assist in planning or structuring a piece of written work. The system modifies its level of response and advice based on the progress of the student through the discussion and on their ability to demonstrate understanding of an area.

Figure 2 TeachBot Interface

Building on the native functionality of the Lingubot system, the system provides extensive logging and analysis capabilities. This allows the learner's progress to be monitored, through analysing their responses and their path through the various building blocks in which the subject matter is organised. Students are able to direct the TeachBot to the area where they want support – if this is in the middle of a topic discussion then their current state is recorded to enable them to return subsequently. The TeachBot can guide students to identified areas of weakness, but it will not prevent them from revising a topic at which they have already shown themselves proficient, if the student wishes.

Finally, the system includes feedback to teachers on the progress of students, both individually and as a group. It can also offer real-time alerts to a teacher via SMS or email, allowing the teacher to provide additional support where necessary.

The English TeachBot system has recently successfully completed trials in secondary schools in Bristol. Some seventy Year 9 (aged 13-14) students used TeachBot to support them while they undertook a non-fiction writing task. In post trial questionnaires, students reported that the system was helpful, easy to use, and fun. Teachers reported a much greater degree of engagement with the tasks than they would have expected, and that it had built confidence in the students, reminding them of material that they had been taught and helping them to apply it. The system is currently undergoing further developments to expand the content and improve performance, prior to large scale trials aimed at assessing the learning gain achieved through using the system.

5 Questions Facing Conversational Agent Developers

There are a wide range of issues and challenges faced by developers and researchers using conversational agents in e-learning systems. Section 5.1 considers the conversation process itself, raising issues of conversation scope, control and structure. Section 5.2 addresses selected technical issues relating to the implementation of conversational technology, which can have far reaching effects on system architecture.

5.1 *Conversational Process – Scope, Control and Structure*

What is the conversational agent going to offer; what is its scope? The scope has substantial implications on system design. While conversational agents have demonstrated e-learning benefits, decisions must be taken about the role of the agent (is it a peer, instructor, motivator, critic?) and about its level of expertise. It must be established what the conversational agent will provide – tutoring, discussion, negotiation, information presentation, answering help queries, companion-

ship, or another function? This establishing of agent purpose informs further design, including:

Who controls the direction of the interaction - user or conversational agent? This question relates to conversation interaction, including initiative (can both the user and the agent introduce conversation topics?), and direction of information flow (does the system present content regardless of user input, or do user inputs influence the outputs of the conversational agent?). For example, a system designed as an interface to an FAQ system may be largely information presentation (based on user requests), while a design allowing greater user proactivity will be more complex in order to ensure that these inputs can be dealt with appropriately. Agent proactivity may also necessitate access to external systems' data if this is to be incorporated in the conversational agent's utterances. If the conversational agent is to be proactive in initiating conversation, then this may need to be constrained to certain circumstances, for example when the user is not busy elsewhere in the system.

What freedoms do users have through the interface; what must you prevent? In the case of the learner model described in section 4, the intention is that user responses may lead to modifying data held in the learner model. In the case of an FAQ system, you would not expect user inputs to affect either the system's conversational knowledge base, or its external FAQ structure. More complex might be a system where the user is allowed to make some modifications, but not others, or perhaps may only make changes in certain circumstances, again demonstrated in section 4.

Do you need to structure the paths along which learners may proceed? This relates to the purpose of the system. Are users expected to focus on a particular task (and perhaps be prevented from straying), or are they intended to explore the conversational dimensions as they choose? In a very focused task user inputs may need to be closely guided, for example by the use of modal dialogue boxes or restricted multiple choice options. Some systems value allowing the learner to develop a social interaction with the conversational agent, and so may allow some off-topic chat. The examples in section 4 support this 'smalltalk' capability to assist in building rapport with the conversational agent, but employ counters to ensure that the user does not remain in smalltalk for too long before being guided back to the learning task.

What will users say? This relates to the knowledge elicitation which may be required early in development, soon after the scope of the conversational agent is determined. While dialogue systems must be able to handle unexpected inputs, understanding what inputs are expected is crucial. Indications of likely content may be provided by techniques such as Wizard-of-Oz studies (where a human 'wizard' simulates part of the behaviour of a final system), by analysis of existing FAQ and help files, or through early prototyping with representative users.

How do you handle unexpected responses elegantly? All natural language systems will face unexpected inputs at some point, either through the user employing language that hadn't been anticipated, or through typing, spelling or grammatical

errors. A certain amount of tolerance to errors can be achieved through simple techniques such as spell checking, but a strategy is required to ensure that users receive suitable responses, even if their input was not successfully or completely parsed. At a basic level this could be by asking the user to correct or re-enter their input, possibly telling the user which part of their utterance was understood, and requesting that they rephrase the unrecognized element. [26] employs a more complex solution, which specifically asks the user to define unknown concepts before adding them to its knowledge base and mapping their semantic relation to existing concepts. Underpinning these techniques, most systems employ a simple "safety net", which catches errors not handled by other techniques and produces a response that ensures that the conversation proceeds.

How do you handle synonymous expressions? Just as there will always be unexpected inputs, even expected responses predicted by developers may be expressed in multiple ways. Simple pattern matching is unlikely to be sufficient for a robust conversational system, as it could result in phrases going unrecognised while synonyms are successfully handled. Similarly, a statement may be made with alternative grammar or voices, and each should be recognised. In addition to synonyms, local argots and Internet or text slang may need to be incorporated; the prevalence of this in Internet chat may influence the language used by learners in a conversational learning scenario.

5.2 *Technical Issues – Implementation*

What approach to natural language understanding should you adopt? There are a range of Natural Language Understanding (NLU) techniques that can be applied in the development of conversational agents. They can be categorised into two main groups: word/phrase matching and softer statistical techniques. Word/phrase matching techniques, as used in the example systems above, attempt to match the user's input with specific words or phrases, in order to categorise a user's response. Practical systems, such as those based on AIML (www.alicebot.org) offer the ability to encode synonyms, to offer flexibility in expected responses. More capable systems (including Lingubot) offer powerful macro capabilities that allow input recognition based on arbitrarily complex sentence structures as well as the constituent words and synonyms. The challenge of establishing a wide and rich enough set of recognitions to cope with the potential range of user inputs (and expected misspellings) should not be underestimated. For example in CALMsystem, the recognition for "belief" includes over 100 alternative words and phrases. This task is simplified to some extent through the use of predefined word lists available in some packages, but even these need to be checked and amended to support the specialist use of language in the subject domain.

Statistical NLU techniques generally rely on an analysis of word frequency and associations across a wide corpus of appropriate texts. Latent Semantic Analysis

(which is used, for instance, in AutoTutor [1]) is a typical example of this type of approach, and can provide an estimate of similarity between text fragments. This in turn allows user inputs to be matched to expected inputs, and appropriate responses generated. Such techniques can be very powerful, but their development and application is more complex than a typical word matching based approach.

How do you select software to implement your selected NLU technique? Once an NLU approach has been selected, a further decision has to be made regarding the software to be used. For word based approaches, both open source and commercial solutions are available. A notable example of open source natural language recognition software is the successful ALICEbot family of systems. ALICEbot implementations are available in a range of programming languages, and are based on a based on AIML, an XML compliant language for programming responses and recognitions. AIML includes extensions allowing the software to invoke calls to other systems, extending the capability of the system beyond simple recognition and response behaviour, and allowing more complex behaviour, including the querying of databases or additional forms of language processing. Commercial natural language engines (such as Lingubot) offer the developer a pre-developed NLU toolset and generally provide a wider range of facilities (such as logging and conversation analysis), as well as product support. As an alternative to these solutions, developers may wish to code their own systems.

How much testing will be required? Conversational interfaces often require considerable testing and iteration cycles before they can be considered to be mature enough for unsupervised use [29]. Commercial systems often undergo over two thousand conversations, and associated updates and corrections, before reaching a fully operational state. This testing can have a significant effect on project timescales.

6 Summary

This paper has explored the variety of conversational agent applications and techniques in the e-learning literature, identifying a variety of agent purposes and strategies including tutoring, language practice, learning companions, pedagogical agents, question answering and encouraging learner reflection. Reported benefits to learners include improvements in grades, motivation, engagement and metacognitive skills. Teachers may also benefit from the ability of conversational systems to provide assessment, reporting and additional classroom tools. Through detailed examples we have described two innovative uses for conversational agents in e-learning, and demonstrated their use as tools in a larger e-learning system and for the provision of support to parallel classroom activities. We have explored a wide range of issues relating to the development of conversational agents, including questions regarding the design of conversational process and issues relating to technical implementation. We conclude that conversational agents have a valuable

role to play in future e-learning and blended learning systems and we expect their use to become increasingly common and progressively more capable as this technology continues to develop.

References

1. Graesser, A.C., N.K. Person, and D. Harter, *Teaching Tactics and Dialog in AutoTutor.* International Journal of Artificial Intelligence in Education, 2001. **12**: p. 23-39.
2. Heffernan, N.T., *Web-Based Evaluations Showing both Cognitive and Motivational Benefits of the Ms. Lindquist Tutor*, in *Artificial Intelligence in Education*. 2003, IOS Press: Amsterdam. p. 115-122.
3. Aleven, V., K. Koedinger, and K. Cross, *Tutoring Answer Explanation Fosters Learning with Understanding*, in *Artificial Intelligence in Education*. 1999, IOS Press: Amsterdam. p. 199-206.
4. De Pietro, O. and G. Frontera, *TutorBot: an application AIML based for Web-Learning.* Advanced Technology for Learning, 2005. **2**(1): p. 29-34.
5. Feng, D., E. Shaw, J. Kim, and E. Hovy, *An Intelligent Discussion-Bot for Answering Student Queries in Threaded Discussions*, in *2006 International Conference on Intelligent User Interfaces*. 2006, ACM Press: Sydney, Australia. p. 171-177.
6. Taylor, K. and S. Moore, *Adding Question Answering to an E-Tutor for Programming Languages*, in *AI-2006, 26th SGAI International Conference on Innovative Techniques and Applications of Artificial Intelligence*. 2006, Springer: Cambridge, UK. p. 193-206.
7. Fryer, L. and R. Carpenter, *Bots as Language Learning Tools.* Language Learning and Technology., 2006. **10**(3): p. 8-14.
8. Abu Shawar, B. and E. Atwell, *Fostering language learner autonomy via adaptive conversation tutors*, in *Corpus Linguistics*. 2007: Birmingham, UK.
9. 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.
10. 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.
11. Baylor, A.L. and Y. Kim, *Simulating Instructional Roles through Pedagogical Agents.* International Journal of Artificial Intelligence in Education, 2005. **15**(2): p. 95-115.
12. Dillenbourg, P. and J. Self, *People Power: A Human-Computer Collaborative Learning System* in *Intelligent Tutoring Systems*. 1992, Springer-Verlag: Berlin-Heidelberg. p. 651-660.
13. Grigoriadou, M., G. Tsaganou, and T. Cavoura, *Dialogue-Based Reflective System for Historical Text Comprehension*, in *Workshop on Learner Modelling for Reflection at Artificial Intelligence in Education*. 2003.
14. Kerly, A., R. Ellis, and S. Bull, *CALMsystem: A Conversational Agent for Learner Modelling.* Knowledge Based Systems, 2008. **21**(3): p. 238-246.

15. Litman, D.J. and S. Silliman, *ITSPOKE: An Intelligent Tutoring Spoken Dialogue System*, in *Human Language Technology Conference: North American Chapter of the Association for Computational Linguistics*. 2004: Boston, MA.
16. Aleven, V., A. Ogan, O. Popescu, C. Torrey, and K. Koedinger, *Evaluating the Effectiveness of a Tutorial Dialogue System for Self-Explanation*, in *Intelligent Tutoring Systems*. 2004, Springer-Verlag: Berlin Heidelberg. p. 443-454.
17. Chou, C.-Y., T.-W. Chan, and C.-J. Lin, *Redefining the Learning Companion: the Past, Present and Future of Educational Agents.* Computers & Education, 2003. **40**: p. 255-269.
18. Rickel, J. and W.L. Johnson, *Animated agents for procedural training in virutal reality: Perception, cognition, and motor control.* Applied Artificial Intelligence, 1999. **13**: p. 343-382.
19. Chan, T.-W. and A.B. Baskin, *"Studying with the Prince": The Computer as a Learning Companion*, in *Intelligent Tutoring Systems*. 1988. p. 194-200.
20. Aimeur, E., H. Dufort, D. Leibu, and C. Frasson, *Some Justifications for the Learning By Disturbing Strategy*, in *Artificial Intelligence in Education*. 1997, IOS Press: Amsterdam. p. 119-126.
21. Kim, Y., *Desirable Characteristics of Learning Companions.* International Journal of Artificial Intelligence in Education, 2007. **17**(4): p. 371-388.
22. Creative Virtual. *UK Lingubot Provider, Major Customers List*. 2004-2008 [Accessed 27/03/08]; Available from: www.creativevirtual.com/customers.php.
23. Jia, J., *The Study of the Application of a Keywords-Based Chatbot System on the Teaching of Foreign Languages*. 2002, University of Augsburg. .
24. Kerly, A. and S. Bull, *Children's Interactions with Inspectable and Negotiated Learner Models*, in *Intelligent Tutoring Systems*. In press, Springer-Verlag: Berlin Heidelberg.
25. Beun, R.-J., E. de Vos, and C. Witteman, *Embodied Conversational Agents: Effects on Memory Performance and Anthropomorphisation*, in *Intelligent Virtual Agents*. 2003, Springer-Verlag: Berlin Heidelberg. p. 315-319.
26. 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.
27. Dimitrova, V., *STyLE-OLM: Interactive Open Learner Modelling.* International Journal of Artificial Intelligence in Education, 2003. **13**: p. 35-78.
28. Kerly, A., P. Hall, and S. Bull, *Bringing Chatbots into Education: Towards Natural Language Negotiation of Open Learner Models.* Knowledge Based Systems 2007. **20**(2): p. 177-185.
29. Spärck-Jones, K., *Natural language processing: a historical review*, in *Current Issues in Computational Linguistics: in Honour of Don Walker*, A. Zampolli, N. Calzolari, and M. Palmer, Editors. 1994, Kluwer: Amsterdam. p. 3-16.