Available online at www.sciencedirect.com
ScienceDirect
Knowledge-Based Systems 20 (2007) 177–185
Knowledge-Based SYSTEMS
www.elsevier.com/locate/knosys

Bringing chatbots into education: Towards natural language negotiation of open learner models

Alice Kerly [a,*], Phil Hall [b], Susan Bull [a]

[a] Electronic, Electrical and Computer Engineering, University of Birmingham, Edgbaston, Birmingham B15 2TT, UK
[b] Elzware Ltd., 70 Mendip Road, Windmill Hill, Bristol BS3 4NY, UK

Received 10 October 2006; accepted 16 November 2006
Available online 11 December 2006

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.

Keywords: Conversational agents; Chatbots; Learner modelling; Negotiation; Natural language

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

* Corresponding author.
E-mail addresses: alk584@bham.ac.uk (A. Kerly), phil.hall@elzware.com (P. Hall), s.bull@bham.ac.uk (S. Bull).

0950-7051/$ - see front matter © 2006 Elsevier B.V. All rights reserved.
doi:10.1016/j.knosys.2006.11.014

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. [17–19]), 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. [20–24]).

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 and 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