AC 2011-1631: WORK IN PROGRESS: THE STUDY BUDDY, A VIRTUAL TUTORIAL AGENT

Keith Garfield, Embry-Riddle Aeronautical University

Dr. Garfield is an Assistant Professor in the department of Electrical, Computer, Software, and Systems Engineering at Embry-Riddle Aeronautical University. His research interests include the use of virtual intelligent agents in educational and training settings.

Work in Progress: The Study Buddy, a Virtual Tutorial Agent

Introduction:

This paper describes an interactive tutorial agent, named the Study Buddy that uses a natural language interface to accompany the student in unguided exploration and discovery of course topics. The Study Buddy is capable of providing concept definitions, provide examples, respond to student questions, query student understanding, and recommend additional web materials for reference. The Study Buddy is novel in that it allows students to create their own examples to be used in the tutorial session. While this study focused on content related to a single course, the techniques used to define and direct agent behaviors is intended to be adaptable to new knowledge domains. This effort is compatible with a general trend in computer use to develop natural human-computer forms of interaction, and also with trends in computer based training to develop virtual coaches and tutors.

Today's students have access to a bewildering amount of data brought to them via an astounding array of computational and communication technologies. There is a growing array of online educational tools exhibiting a wide array of functionality. These tools range from simple replication of course content, to additional course content, to virtual labs, virtual environments[2], and interactive virtual agents acting as coaches or tutors [5]. Content may be static, with fixed sequences, or exploratory, allowing students to discover content on their own.

Unfortunately, this unprecedented access to information does not necessarily result in more effective learning. Ideally, students would explore this wealth of information independently, but their explorations are inhibited by the lack of structure, time, and motivation. Additionally, the limited interactions provided by most web technologies offer little help to students who require an alternative, as opposed to repetitive, presentation of the material.

The Study Buddy attempts to overcome these obstacles by providing an engaging, natural language conversation with the student during the tutorial session. The tutorial sessions are guided by a set of agent behaviors that map to specific course learning outcomes. This reliance on learning outcomes is used to overcome obstacles encountered in natural language interactions and provide reasonable reactions to student inquiries within the target subjects, while redirecting the conversation when students are off topic. The Study Buddy maintains a model of student understanding to guide the tutorial session if the student does not wish to, or know how to, frame specific questions.

The agent has been developed and tested for functionality, but is not yet deployed. Thus, no field test data was acquired to assess the agent's effectiveness in terms of motivating students to explore classroom topics outside of the class. The work has been beneficial in establishing the feasibility of the approach, identifying an architecture that may be used for more general course topics, and developing guidelines for effective use.

Goals:

The larger goal of the Study Buddy project is to provide a non-threatening, engaging tool to encourage the student to explore course topics outside of the classroom. Specific objectives for this project are to develop and field a prototype Study Buddy agent, assess its effectiveness in terms of motivating students, and utilize the work as a basis for future growth in tutorial agent capabilities.

This project targeted a purposefully small and well-defined set of course outcomes for a math course that described properties of certain sets, called relations. The outcomes require the students: 1.) be able to recognize whether or not the relations exhibited reflexive, symmetric, and/or transitive properties, and 2.) be able to create relations that either did or did not exhibit those properties. For example, a typical classroom exercise would be to present an example relation and require the students to state whether a specific property was present. If the property was not present, the students are required to modify the relation so that the requested property was maintained. This small set of outcomes is rich enough to require a complex set of behaviors on the part of the Study Buddy, in that the agent must be able to provide definitions, provide examples, accept example problems from the student, query the student, and assess the correctness of student responses.

One objective for this project was to interact with the agent using natural language. Thus, the agent requires natural language processing (NLP) capabilities. While NLP has not been fully solved in a general sense, there are numerous examples of systems demonstrating some success in limited domains.

Since the Study buddy is intended to emulate natural language conversations that often do not flow linearly, the ability to present material in an ad hoc fashion was a goal. A large amount of online instructional material is presented in linear fashion, requiring students to move from the first topic to the last in succession. Virtual exploratory environments are an exception to this format, allowing students to simply explore an environment and encounter material in a non-linear fashion.

To date these goals have been partially met. A prototype agent demonstrating basic conversational capabilities has been developed. The agent was not fielded, and therefore no data was collected regarding the effectiveness of the agent.

Functional Design Description:

The Study Buddy's functional design was driven by the course learning objectives. The objectives effectively imposed clear design requirements in terms of agent behavior and NLP capabilities. The specific course outcomes were selected for the clarity of measurement. As Figure 1 implies, this translated relatively easily into program functions, standardized commands, and NLP requirements to recognize those commands. For example, the course outcome "The student shall be able to identify reflexive relations." translated easily into a program function that presents an example relation to the student and differentiates between variations on "Yes." and "No." responses.

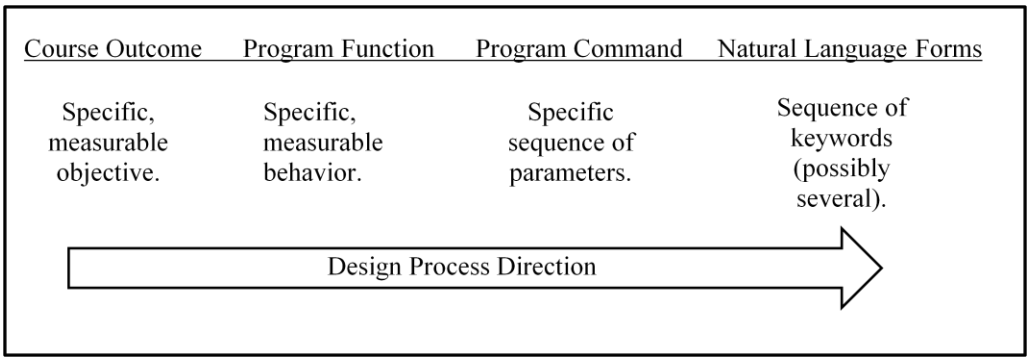

Figure 1: Design Process, Outcomes Drive Behaviors and Language

The set of course outcomes also formed the basis of a conversational framework. The Study Buddy is able to interact with a student by recognizing sequences of keywords in the student's text as requests to perform an action. The set of targeted course outcomes thus defines the set of keywords and keyword phrases recognized by the Study Buddy's NLP component. This basic NLP capability is made more natural by recognizing multiple forms of expressing the same request. For example, "What does reflexive mean?", "Define reflexive.", and "What's the definition for reflexive." may all map to the single program function that provides the definition.

Using this framework, the Study buddy is also able to track which topics the student has discussed, and also which topics the student has been assessed on. This allows for the Study Buddy to suggest continued directions for the conversation.

The end result is a chat-like system capable of a set of behaviors sufficient to present and assess the targeted course outcomes. These behaviors, summarized in Table 1, allow reasonable reactions to student inquiries within the target subjects, and redirect conversation when students are off topic. The remainder of this section discusses design decisions tied to specific design challenges or goals.

Natural language understanding and interaction have been widely studied in a variety of domains[1, 7]. While the problem has not been solved in the general sense, advances have been made in limited domains. The main issues associated with automated understanding of natural language understanding can be summarized as relating to complexity of knowledge and ambiguity of speech. The natural language recognition problem is resolved using a keyword approach to parse input, as opposed to attempting to represent a deep semantic understanding of the sentence. Previous work has shown success using this approach in limited topic domains[3, 4]. Using this technique, keywords and their placement in a sentence are mapped to agent behaviors. As shown in Figure 2, sentences are mapped first into an intermediate representation where the keywords are identified and tagged as to grammatical function, and then again mapped to standardized commands that initiate Study Buddy behaviors. Each command requires a complete set of parameters (such as action required, set under discussion, and relation under discussion) for an action to be initiated. Incomplete commands are completed by querying the student ("Which relation are you talking about?"). Once a complete command is formed, the Study Buddy performs the desired action. Study Buddy natural language responses are fairly scripted, with variations selected randomly to avoid overt repetition.

Table 1: Summarized Capabilities of Study buddy

Capability	Description
Provide Formal Definitions	Upon request, provide a formal definition of the Reflexive, Symmetric, and Transitive properties of relations.
Provide Informal Definitions	Upon request, provide an informal definition of the Reflexive, Symmetric, and Transitive properties of relations.
Identify Sets	Recall by name pre-defined mathematical sets, or sets created by the student during a session.
Identify Relations on Sets	Recall by name pre-defined relations, or relations created by the student during a session.
Provide Examples	The Study Buddy can generate and display examples of relations in response to student queries.
Query Students	The Study Buddy is able to pose questions to students and assess their responses in terms of the properties listed above.
Verify Student Answers	The Study Buddy is able to recognize correct answers to queries, and provide corrections to incorrect answers.

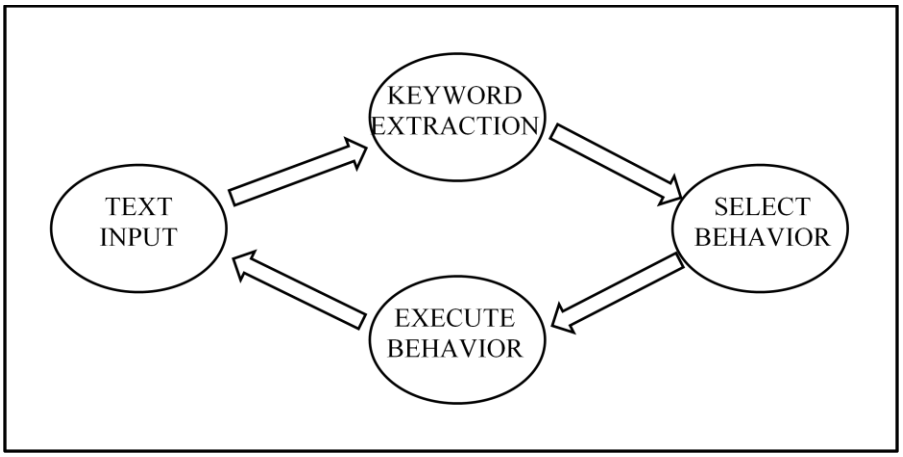

Figure 2: Interaction Cycle

The Study Buddy is intended to be exploratory in nature. The issues associated with this goal are recognizing what the student intends to do, and providing cues to the hesitant or confused student. This provides a challenge in that material must be accessed in a non-linear fashion. An additional challenge is that the Study Buddy must recognize when a student is expressing an interest in a topic, and when a student has mastered an outcome and may safely move on. In order to allow students to visit topics in an exploratory fashion, the Study Buddy maintains a list of completed course outcomes and compares this list to the entire set of targeted outcomes. By not requiring that these outcomes be satisfied sequentially, the student is able to suggest which topics to visit. Thus, the conversation can unfold in a non-linear fashion as natural conversations do. If the student is unwilling or unable to make suggestions, the Study Buddy uses the list of

outcomes to suggest topics. This is done by maintaining an effective queue of behaviors, with the behavior requested by the student having priority.

Assessment of educational outcomes is typically done in an automated setting through question and response techniques. The non-interactive nature of online, automated sessions makes assessment difficult. If assessments are repeated it is difficult to disambiguate true learning from trial-and-error approaches, and if assessments are not repeated learning may not be measured. The Study Buddy attempts to resolve this by creating variations on examples and allowing the student to create their own example content to work with. The specific course, Discrete Structures, and course outcomes targeted for this study were selected largely due to their unambiguous nature, which assists in discussion and assessment. The Study Buddy contains a set of designed behaviors to assess these outcomes. Students are able to demonstrate comprehension of set properties by classifying an example set as having, or not having, the property. Students are able to demonstrate application of their knowledge by creating sets that have a requested property. In a variation of this, the Study Buddy presents sets that do not contain the property being discussed, and requires that the student modify the relation to obtain the property.

The goal of allowing students to create their own example relations is intended to add to the exploratory and engaging nature of the Study Buddy. In a general sense, the allowance of student content raises the issue of being able to address that specific content in the tutorial session. The system must be able to store the content, recognize when it is being addressed, and recall it. In addition, the tutorial must be able to manipulate the student generated content and query the student as appropriate. For the Study Buddy, these issues are greatly simplified, due to the nature of the behaviors targeted. While the Study Buddy contains pre-defined relations to work with, the student is able to define their own, providing both a name and set elements, for further discussion. The Study Buddy software simply maintains a list of all sets, predefined and created, and is able to treat them interchangeably. Since the Study buddy has no log-in feature, this content is not saved between sessions.

Every educator understands that this challenge is not limited to online content; however, the absence of personal interaction severely limits the capability of the system to engage the student. Specifically, the unnatural mechanisms that may be required to navigate a virtual world may distract or annoy students [6]. People are intuitively aware of this online limitation when utilizing email or engaging in online discussions. The interactive nature, natural language format, exploratory environment, and ability to create individual content are all intended to enhance student engagement. Since the Study buddy has not been deployed, this aspect has not yet been assessed.

Using the techniques outlined above, the Study Buddy is be capable of providing concept definitions, provide examples, respond to student queries about examples, and query student understanding. The focus of the Study Buddy is on engagement and exploration, as opposed to direct measurement and assessment. Assessment of understanding is performed in the context of enhancing and guiding the exploration, as opposed to measuring the success or failure of a student's performance.

Future Work:

More time is required to develop the agent to the point where it may be usefully fielded and tested for engagement and effectiveness in enhancing student learning. Longer term studies are needed to investigate the ability of the techniques used in the Study Buddy to support other educational outcomes, and the difficulty of adapting Study Buddy behaviors to those outcomes.

Conclusions:

A prototype agent demonstrating basic conversational capabilities has been developed. The Study Buddy agent is capable of performing in a simple questions and answer format on the target course topics. The primary conclusion for this work is that the concept of a Study Buddy is feasible. The effect of this agent on student learning has not been tested.

The method of allowing course learning objectives to drive the functional design requirements resulted in clear design goals for the Study Buddy in terms of functionality and NLP needs. Though not field tested, the agent has demonstrated a capability to exhibit the desired behaviors and interact using keyword recognition language processing techniques. While the techniques discussed in this paper were presented in the context of the specific course learning objectives used during the project, it is hoped that other developers of tutorial agents may benefit from this approach.

Bibliography:

1. Allen, James. *Natural language understanding (2nd ed.)*, Benjamin-Cummings Publishing Co., Inc., Redwood City, CA, 1995

2. Dickey, M.D. (2003). 3D Virtual Worlds: An Emerging Technology for Traditional and Distance Learning. Proceedings of Ohio Learning Network, Columbus, OH, March 2003.

3. Garfield, K., Washburn, D., Gerber, M., Kiriwas, A. P. Breaking the Natural Language Barrier in human-Machine Interaction. Interservice/Industry Training, Simulation, and Education Conference. December, 2003, Orlando, FL.

4. Garfield, K., Washburn, D. Natural Language Interaction Between Live and Synthetic Participants in a MOUT Environment. Simulation Interoperability Workshop, September 2002, Orlando, FL.

5 . Graesser Arthur C. ; Shulan Lu ; Jackson George Tanner ; Mitchell Heather Hite ; Ventura Mathew ; Olney Andrew ; Louwerse Max M., *Autotutor: A tutor with dialogue in natural language*. Behavior research methods, instruments & computers. Psychonomic Society, Austin, TX , 2004.

6. Taylor, W. Student Responses to Their Immersion in a Virtual Environment. In *Proceedings of Annual Meeting of the American Educational Research Association*, pp. 12, 1997.

7. Winograd, T. *Understanding Natural Language*. New York Academic Press, 1972.