Targeted Help and Dialogue about Plans

Beth Ann Hockey, John Dowding, Gregory Aist and Jim Hieronymus
Research Institute for Advanced Computer Science
NASA Ames Research Center
Moffett Field, CA 94035
{bahockey,jdowding,aist,jimh}@riacs.edu

1 Introduction

We will demonstrate two novel capabilities:

- Targeted Help: a component for providing users with embedded tutoring about the dialogue system and;
- Dialogue for discussing future plans with a state of the art planner/scheduler

This prototype controls a mobile robot and various station functions in a simlation of the International Space Station. The components of this system extend work from the RIALIST simulated PSA Dialogue System (Rayner et al., 2000). The demo system features an asynchronous dialogue manager (Boye et al., 2000), combined with the Nuance Speech Recognizer, the Gemini language understanding system (Dowding et al., 1993) , and the Festival Speech Synthesizer. The components are integrated using SRI's Open Agent Architecture (OAA). It includes new agents for performing Targeted Help. In addition the system's language coverage includes temporal expressions for time points and durations, which are crucial to discussions of future plans. The system integrates capabilities for discussing and directing current activities, for queries about past states and for interactive future planning. Targeted Help is available across all interaction types.

1.1 Targeted Help: Guiding Users Toward In-Coverage Utterances

Grammar based recognizers tuned to a domain perform very well, in fact better than comparable Statistical Language Models (SLMs) for in-coverage utterances (Knight et al., 2001). However, users, particularly non-expert users of dialogue systems, will produce utterances that are out of coverage. Targeted Help makes use of these out-of-coverage utterances to produce informative and natural utterances that will guide the user towards producing more in-coverage language. The utterances produced by Targeted Help are composed of one or more of the following pieces:

- Report of backup SLM recognition hypothesis,
- Descripition of the problem with the user's utterance (e.g. missing vocabulary); and
- A similar in-coverage example.

Since we are essentially embedding a very simple instructional system about coverage into the dialogue system our work has points in common with instructional dialogue systems. There are also similarities between this approach and work on call routing.

This is a novel approach to a real problem. Many types of systems will have few repeat users with the result that most users will be inexpert. However, even in systems where training is possible, users can remain inexpert surprizingly long. In a study done on the original RIALIST simulated PSA system, naive users were given a half hour training session using the dialogue system before beginning the series of tasks for the system test. Even after training users produced at least 8.5% out of coverage utterances. In addition to providing users with information about the system that is not explicitly available (e.g. what words or grammatical forms are out of coverage), Targeted Help takes advantage of lexical entrainment by providing similar in-coverage examples. In addition, embedding this training in the system takes advantage of the effectiveness of immediate feedback as an aid to learning.

The Targeted Help component is a separate agent which analyzes the back up recognizer output to produce an informative system response. To implement the Targeted Help capability we also add a second recognizer using a category based SLM running in parallel to the usual grammar based recognizer. This speeds production of the Targeted Help feedback by insuring that the back-up recognition hypothesis is immediately available when the main recognizer fails. The Targeted Help component has 2 subroutines, one that produces a prompt based on a diagnosis of the problem with the user's utterance, e.g. unknown words, unknown grammatical structures, bad endpointing. The other component produces an in-coverage example sentence based on the form of and lexical items recognized in the user's utterance. We match the example to the sentence type of the user's utterance: wh-question, yn-

Proceedings of the ACL-02 Demonstrations Session, Philadelphia, July 2002, pp. 100-101. Association for Computational Linguistics.

question, answer, or command. Furthermore, for commands, which are a large percentage of the utterances we also attempt if possible to produce an in-coverage example utterance containing the same verb and arguements. The recognition hypothesis of the SLM is also reported as part of the system's response to the user.

1.2 Dialogue about Plans

This part of the demonstration will describe a work in progress towards building a spoken dialogue interface to the Extensible Universal Remote Operations Planning Architecture (EUROPA) (Frank et al., 2000) planner for the purposes of planning and scheduling the future activities of the PSA. The spoken dialogue interface gives the user the capability to ask for a description of the plan, ask specific questions about the plan, and update or modify the plan. We anticipate that a spoken dialogue interface to the planner will provide a natural augmentation or alternative to the visualization interface, in situations in which the user needs very targeted information about the plan, in situations where natural language can express complex ideas more concisely than GUI actions, or in situations in which a graphical user interface is not appropriate.

The Mars Pathfinder mission landed a semi-autonomous robot, Sojourner, on the surface of Mars which operated for 85 days, from July 4, 1997 to Sept 27, 1997 (MPF, 1997). At the end of each day, Sojourner would broadcast data back to earth, then hibernate for the Mars night. On Earth, teams of scientists and engineers would engage in marathon planning sessions at night to plan the next day's activities. If human planners were able to correctly and completely specify all the constraints and priorities on a planning problem of this magnitude at the outset, then this would just be a hard planning problem. But, human planners do not work that way. They can come up with an initial estimate of the constraints, but that may not be satisfiable, in which case they are faced with choices of which constraints to consider removing. Even when the set of constraints is satisfiable, when the human planners are given a description of the current plan, they may come up with additional constraints and priorities in an attempt to improve the plan. Thus, constraint and priority generation is a process that proceeds iteratively with planning, and requires increased sophistication in the planner interface.

1.2.1 EUROPA

Researchers at NASA's Ames Research Center and the Jet Propulsion Laboratory, are developing new planning technology to address this problem: EUROPA. EUROPA is a system for representing and reasoning about complex temporal plans involving constraints on time, resources, and routes. EUROPA is the descendent of the Remote Agent Planner (RAP) (Jónsson et al., 2000). RAP was deployed as part of the control system for the Deep Space One spacecraft in May of 1999.

EUROPA manages plans by continuously reformulating the planning problem as a Dynamic Constraint Satisfaction Problem (DCSP). This reformulation is done by mapping each partial plan to a CSP. The temporal constraints form a Simple Temporal Network, which can be efficiently solved, while the rest of the constraints form a general, non-binary CSP represented by procedural constraints (Jónsson and Frank, 2000). Planners may interact with the EUROPA interface by modifying plans, asking whether a plan is consistent, and asking whether any more work must be done to complete the plan.

References

J. Boye, B.A. Hockey, and M. Rayner. 2000. Asynchronous dialogue management: Two case-studies. In *Gotalog: Fourth Workshop on the Semantics and Pragmatics of Dialogue*.

J. Dowding, M. Gawron, D. Appelt, L. Cherny, R. Moore, and D. Moran. 1993. Gemini: A natural language system for spoken language understanding. In *Proceedings of the Thirty-First Annual Meeting of the Association for Computational Linguistics*.

J. Frank, A. Jónsson, and P. Morris. 2000. On reformulating planning as dynamic constraint satisfaction. *Symposium on Abstraction, Reformulation and Approximation*.

A. Jónsson and J. Frank. 2000. A framework for dynamic constraint reasoning using procedural constraints. *Euopean Conference on Artificial Intelligence*.

A. Jónsson, P. Morris, N. Muscettola, K. Rajan, and B. Smith. 2000. Planning in interplanetary space: Theory and practice. In *Proceedings of the Fifth International Conference on Artificial Intelligence Planning and Scheduling*.

S. Knight, G. Gorrell, M. Rayner, R. Koeling, and I. Lewin. 2001. Comparing grammar-based and robust approaches to speech understanding: A case study. In *Proceedings of Eurospeech 2001*, Aalborg, Denmark.

1997. *Mars Pathfinder Fact Sheet*. http://www.jpl.nasa.gov/news/fact_sheets/mpf.pdf.

M. Rayner, B.A. Hockey, and F. James. 2000. A compact architecture for dialogue management based on scripts and meta-outputs. In *Proceedings of ANLP 2000*.