NATURAL LANGUAGE DIALOGUE FOR INTERACTIVE PLANNING

James F. Allen and Len Schubert
Department of Computer Science
University of Rochester
Rochester, NY 14627

PROJECT GOALS

The goal of this project is to develop the underlying technologies for spoken dialogue systems to serve as highly interactive interfaces to AI-based reasoning systems. Most current speech and natural language projects are focusing on applications that involve only limited dialog, and little intelligent reasoning, such as data-base query and form-filling applications. But the great promise for speech and natural language interfaces is in providing useful interfaces to complex reasoning systems such as planning systems and expert systems.

The techniques developed are being incorporated into an integrated dialog system set in a simulated transportation domain (the TRASINS domain). The user interactively develops a plan in a mixed-initiative interaction with the system. The development of the system is supported by a corpus of person-person spoken dialogs collected in the TRAINS domain. This corpus has been used for a wide range of analyses, including the detection of speech repairs, the analysis of discourse structure, and the development of a grammar and parser for interactive spoken dialog. The system currently works from keyboard input based on transcripts, but we are in the process of adding a speech recognizer.

RECENT RESULTS

1) We developed a stochastic technique for detecting and realizing speech repairs in spoken dialog. The technique is designed to work from the words output by a speech recognizer, and the only assumption made beyond current state of the art is the recognition of word fragments. Using stochastic methods with no prosodic information or parser, we are able to detect and correctly realize over 90% of the repairs in our corpus. See paper in the HLT proceeding for details.

2) We have designed a new discourse reasoner that maintains what each agent believes, what each agent has suggested about the plan, and what parts of the plan have been agreed on so far. The module can filter possible speech act interpretations using knowledge of the agents' beliefs and knowledge of the current plan, and works in conjunction with a domain reasoner that performs the plan reasoning tasks in the TRAINS world.

3) We have constructed a grammar (syntax and semantics) that covers over 500 utterances in the TRAINS domain, including several full dialogues. The system can adapt

lexical entries from the Alvey lexicon, providing access to the syntactic features of over 6000 words. We have developed a stochastically-driven chart parser that identifies the correct parse 70% of the time in our preliminary tests.

4) We have implemented a system that, when presented with a word it has never seen before, creates a new lexical entry with meaning postulates that represent a partial semantic definition. The algorithm uses a model of the word formation process (e.g., affixation, argument structure alternations, compounding etc.) to identify the syntactic class an approximate semantic definition.

5) We have collected an additional 8 hours of TRAINS dialogs, and have produced an word-aligned transcription and annotated all the repairs (approx. 1000 instances). This corpus has been used as the basis for our system to detect and realize repairs, and to develop our grammar and train the parser.

PLANS FOR THE COMING YEAR

The most significant goal for the coming year is to add a speech recognition front-end to the TRAINS system. This will allow us to more realistically explore issues in parsing actual dialog. One of the first tasks to face will be to develop some techniques for utterance segmentation. In a dialog, a speaker often makes several separate utterances within a single turn, and it is crucial for the system to recognize the utterance boundaries. We expect to start exploring some simple prosodic cues to utterance boundary locations, as well as using stochastic and syntactic constraints.

We plan to collect additional dialogues in the TRAINS domain, and to continuing the annotation of the existing corpus. This year we intended to do an extensive annotation of the referring phrases in the corpus, and to annotate prosodic features using the ToBi annotation scheme.

A longer term concern is robustness. The TRAINS system can handle some reasonably size dialogues containing twenty or so turns. It is very fragile, however, because there are many different levels of analysis, each covering a slightly different range of phenomena. One of the central tasks for the coming year is to try to identify the design decisions that make the system brittle and to redesign the system in order to improve robustness.