Linguistic Knowledge Sources for Spoken Language Understanding

Deborah A. Dahl

Unisys Center for Advanced Information Technology
Paoli, Pennsylvania 19301

PROJECT GOALS

The objective of the Unisys Spoken Language Systems effort is to develop and demonstrate technology for the understanding of goal-directed spontaneous speech. The Unisys spoken language architecture couples speech recognition systems with the Unisys discourse understanding system, PUNDIT. PUNDIT is a broad-coverage language understanding system used in a variety of message understanding applications and has been extended to handle spoken input. Its power comes from the integration of syntax, semantics, and pragmatics processing, the ability to port rapidly to new task domains, and from an open, modular architecture. PUNDIT is unique in its ability to handle connected discourse; it includes a reference resolution module that tracks discourse entities and distinguishes references to previously mentioned entities from the introduction of new entities. The PUNDIT front-end supports turn-taking dialogue and permits the system to include both questions and answers in building an integrated discourse context, required for the handling of interactive communication. PUNDIT has been interfaced to several speech recognition systems (MIT SUMMIT and ITT VRS-1280) to perform applications in direction-finding assistance, air travel planning and air traffic control.

RECENT RESULTS

We have begun to explore a new approach to comparing speech recognition systems and for investigating speech/language interactions in spoken language systems. This approach involves coupling a single natural language system to multiple speech recognition systems. At this point we have worked with output from MIT, BBN and Lincoln Labs, and are about to start working with output from Dragon. We believe this will be extremely useful for giving us a handle on the tradeoffs between recognizer improvements and improvements in spoken language understanding.

We have also recently made improvements in semantic processing which allow semantics to handle, in a rule-governed way, a variety of syntactic constructions which exhibit non-transparent argument structure. These improvements led to an improvement of our benchmark ATIS score to 48.3% from 36%.

Our pragmatics processing has been enhanced to introduce information presented in displays into the dialog and to support subsequent references by the user to items in the display.

We have also developed a new approach to training the semantic component of the system using hand-generated rules developed for 900 sentences of the ATIS corpus to provide a basis for hypothesizing the argument structure of unknown verbs.

In collaboration with the University of Pennsylvania, we have recently implemented a probabilistic chart parser with an unknown word model and have partially integrated it into the system. This parser differs from previous attempts at stochastic parsers in that it uses a richer form of conditional probabilities based on context to predict likelihood.

PLANS

During the year to come we plan to:

- Evaluate speech/language integration with multiple speech recognizers

- Use Lincoln Labs stack decoder to provide tight speech/language coupling

- Begin development of a new performance task – information retrieval interface

- Extend trainable semantics

- Improve dialogue processing for full dialogues