A Real-Time Spoken-Language System for Interactive Problem-Solving

Patti Price and Robert C. Moore
SRI International

Objective:

SRI is developing a system to improve complex problem-solving through the use of interactive spoken language in conjunction with other media. This requires real-time performance, large vocabulary, high semantic accuracy and habitability, as well as robustness to expected and unexpected variability. SRI's spoken-language system is being developed in the air travel planning domain along two overlapping research and development lines, one focussed on an SLS kernel for database query, and the other on the full interactive system.

Recent Accomplishments

- Improved natural language understanding in the air travel planning domain (about 90% of a test-set of 319 sentences collected at TI now parse, 64% receive a semantic interpretation, over 25% retrieve the correct answer from the database.
- Implemented corrective training to improve recognition performance; on the standard training set this improves speaker-independent perplexity 60 performance from 6.7% error to 5.1% error, and for a larger training set (about 11,000 sentences), improves speaker-independent recognition from 5.3% error to 4.1% error.
- Evaluated the effects on performance of larger training sets: using about 11,000 training sentences compared to about 4000 training sentences improves from 6.7% error to 5.3% error.
- Started initial experiments in consistency modeling: showed significant improvement for separate modeling of male and female speech for both the standard and the large training sets (performance improves from 6.7% error to 5.7% error on the standard training and from 5.3% error to 4.3% error on the large training set). Speaker-independent perplexity 60 performance can be expected to improve still further when gender-specific modeling with corrective training on large data sets is implemented.
- Started implementing procedures for porting to new vocabularies, including the ATIS vocabulary, in-

cluding the automatic generation of baseforms, applications of rules and the creation of word-models based on existing training data from other domains (currently about 9000 verified words are in the dictionary and another 20,000 from the Brown corpus most frequent words in process).

- Incorporated in the SRI Speech Understanding System techniques (developed under an NSF contract) that use statistical information about the speech signal to improve recognition accuracy in noisy environments; started the real-time implementation of these algorithms.
- Started port of phonology software from LISP to C, and demonstrated performance improvement through the use of more detailed, statistically trained phonological models.
- With help from TI, implemented a functional equivalent of the TI data-collection mechanism; collected and transcribed data from 10 subjects using the TI protocols. Designed experiments for variations on this mechanism and started data collection and analysis.

Plans

- Continue ATIS development.
- Expand effort in consistency modeling and training on larger data sets.
- Continue development of software for easy porting to new vocabularies.
- Begin effort to detect out-of-vocabulary items and design and evaluate an appropriate interface.
- Collect and analyze more human-machine data under various conditions in order to guide the design and evaluation of the SRI SLS.