School of Computer Science

Speech Consortium

School of Computer Science

Carnegie Mellon's Computer Science Program has established an international reputation as a leading center for computer science education and research. Research is organized into four primary areas: programming systems, computer systems, artificial intelligence, and theoretical computer science. The School includes approximately 80 faculty, 115 technical and administrative staff members, 175 graduate students, and 10 visiting scientists.

Numerous governmental and industrial organizations—including the DoD, NSF, NIH, IBM, Digital Equipment, Intel, Sun, MicroSoft and Apple Computer—generously support our work.

Speech Research Laboratory

Since its inception over twenty years ago, the Carnegie Mellon Speech Research Laboratory has made major contributions to the field. In 1972 we developed the first connected-speech understanding system, HEARSAY-I. Work on the DRAGON system (1975) led to formulating the speech interpretation problem as one of hidden Markov modeling. HEARSAY-II (1976) introduced the "blackboard model" as a problem-solving mechanism for systems having diverse, cooperating knowledge sources.

Our HARPY system (1976) was the first to use beam search for solving the exponential graph-search problem. The system was subsequently implemented as a special-purpose architecture with real-time performance: the Harpy Machine (1979).

In 1982 we developed FEATURE, a speaker-independent alphabet recognition system based on phonetic feature-extraction. In 1987 we demonstrated SPHINX—the first 1000-word, speaker-independent, connected-speech recognition system performing at better than 95% word accuracy. Based on SPHINX-II and SPHINX-III, our current systems, such as ATIS, consistently lead the field in performance.

Speech Consortium

To broaden the Speech Laboratory's participation base and to facilitate rapid transfer of new technologies into industry, the School has established the Carnegie Mellon Speech Consortium. The Consortium provides a comprehensive window into speech research, enabling members to obtain state-of-the-art information and to develop technical expertise in automated speech recognition.

Opportunity

Corporations interested in applying speech technology within their operations or products are invited to join the Carnegie Mellon Speech Consortium. Members receive the following direct benefits:

- Eligibility to receive all non-proprietary data and software on a non-exclusive, royalty-free basis.

- The Portable Speech Library: software that can aid a participant to quickly establish a state-of-the-art facility for speech recognition. The Library includes both discrete and semi-continuous recognition systems, plus training software for acoustic and language models and for spoken language parsing (the PHOENIX system). We will update the Library regularly and encourage members to submit enhancements.

- Early notification of new developments and access to source code, to maximize opportunities for commercial advantage.

- Direct contact with established leaders in speech recognition and opportunities to consult with members of the Speech Research Laboratory.

- Early acquaintance with graduate students who will lead future research in the field.

Corporate membership in the Speech Consortium is based on an initial contribution of $250,000, renewable annually at $125,000.

It is possible to arrange for employees of the Consortium member to join the Speech Research Laboratory for a period of up to a year, at a cost of $65,000 per employee. Doing so creates the opportunity to train individuals and enhances the possibilities for technology transfer.

Please contact:

Dr. Alexander I. Rudnicky
Director, Speech Consortium
School of Computer Science
Carnegie Mellon University
5000 Forbes Avenue
Pittsburgh, Pennsylvania 15213
[412] 268-2622
[412] 268-5576 (fax)
air@cs.cmu.edu