A SNAPSHOT OF TWO DARPA SPEECH AND NATURAL LANGUAGE PROGRAMS

Charles L. Wayne

DARPA/ISTO

DARPA is investing in speech and natural language processing research to ensure the availability of key technology needed by the Department of Defense for a wide variety of applications. The research programs aim (a) to develop enabling component technology that can be integrated on demand and/or rapidly tailored for specific applications and (b) to demonstrate that technology in limited prototypes. The programs are highly synergistic and emphasize objective performance evaluation.

This note describes the overall programs; the following project summaries provide additional detail.

SPOKEN LANGUAGE

The DARPA Spoken Language program has two major components: large vocabulary speech recognition, which has many applications, and spoken language understanding, aimed at interactive problem solving. Both deal with spontaneous, goal-directed, natural language speech. And both aim for real-time, speaker-independent or speaker-adaptive operation. The program also includes basic research to fuel the next generation of advances.

Performance evaluation for speech recognition is currently being conducted using the Resource Management (RM) corpus, which consists of read queries and commands, and the Air Travel Information System (ATIS) corpus, which consists of spontaneous queries and commands. Plans are underway to expand the ATIS corpus and to replace the RM corpus with a more challenging one.

Performance evaluation for speech understanding is being conducted with the ATIS corpus, collected from subjects interacting with a simulated (wizard-based) understanding system that contains certain data from the Official Airline Guide (OAG).

In addition, several groups are also developing spoken language technology demonstration applications. The most advanced of these is MIT's Voyager system, which provides navigational assistance for Cambridge, Massachusetts.

Groups currently being funded include BBN, Brown, BU, CMU, Dragon, Lincoln, MIT, SRI, TI, and UNISYS. The program is greatly enriched by the voluntary participation of AT&T in the periodic performance evaluations.

WRITTEN LANGUAGE

The Written Language program is developing the technology needed for large-scale text processing. The program encompasses message understanding, natural language learning, basic research, and corpus building. It will soon include work on machine translation.

Performance evaluation of message understanding systems is done in terms of database template filling. Multisite evaluations take place in message understanding conferences (MUCs). MUC-2, which was held in 1989 used Navy OPREP messages. MUC-3, which is happening in two phases this year, is using FBIS news reports. Performance evaluation of natural language learning techniques also takes place (in part) in the context of the MUC process.

Performance evaluation of machine translation algorithms will also be done on previously unseen, naturally occurring texts. DARPA's MT work is just beginning this year, and an important part of the initial phase will be to develop specific evaluation methodologies.

Groups currently being funded include BBN, Columbia, NMSU, NYU, Penn, Rochester, SRI, and UCB. The program is greatly enriched by the participation of many other groups in the DARPA speech and natural language workshops and in the MUC process.

PROGRESS ON CONTINUOUS SPEECH RECOGNITION

RESOURCE MANAGEMENT CORPUS, READ SPEECH
1000-WORD VOCABULARY, PERPLEXITY 60

EVALUATION —> PROGRESS