Session 2: Natural Language I

Damaris M. Ayuso, Chair
BBN Systems and Technologies Corporation
Cambridge, MA 02138

This first session on natural language covered a wide variety of topics, from formal treatments of semantics to research in multi-media interactions.

The first paper, by L. Schubert and C. Hwang, concentrated on compositionally defining the mapping between a sentence and an appropriate meaning representation that takes tense and aspect into account. In order to derive non-indexical formulae from indexical ones as a function of context, they make use of a novel type of context component, a *tense tree* which evolves throughout the discourse. Occurrences of tense-aspect operators are represented as tree branches corresponding to episode (event) relationships. An interesting feature of these trees is that one can "read off" the tree the relationships implicit in the tense-aspect operators and surface ordering of sentences.

K. McKeown from Columbia University gave a multimedia presentation showing examples of COMET's capabilities. The presentation focused on the recent advances in COMET, in particular, improvements in media coordination, text generation, and graphics generation. In experiments where both pictures and text are generated in an explanation, users were found to strongly prefer that sentence breaks correspond to picture breaks. Consequently the text and graphics generators are now coordinated to ensure appropriate sentence breaks where possible by influencing lexical choice and syntactic structure. In text generation, a framework for lexical choice was developed using the Functional Unification Formalism, which allows the integration of various types of constraints uniformly. Finally, the generation of graphics was made context-dependent in certain areas, e.g., constraints from previous pictures apply to the current picture in order to prevent disconcerting, unnecessary changes in viewpoint.

Although the Y. Schabes and A. Joshi paper covered two recent advances in TAGs, synchronous TAGs and LR-style parsers for TAGs, Schabes focused in his talk on the use of synchronous tags for interpretation. By representing semantic expressions as TAG trees, simple mappings can be defined between trees in the grammar and these semantic trees so that via the synchronous TAG process, a logical form can be derived as a parse tree is built. This approach is similar to the classic syntax-rule/semantic-rule compositional mappings of other NL systems. However, Schabes pointed out that the advantages offered by the extended domain of locality of TAGs in syntax are also useful for semantic purposes, e.g., in treating idioms and long-distance WH-dependencies. The use of synchronous TAGs for machine translation was also mentioned, where TAG trees in the grammars of two languages are related synchronously.

E. Hovy from ISI/USC reported on a new experiment based on an old idea: using a KL-ONE-like knowledge representation system to perform tasks now normally done via unification in unification-based NL systems. The idea is that a single knowledge base should encode both syntactic and semantic knowledge, so that via a single operation, classification, the system will deduce the correct syntactic/semantic description of the sentence, given an initial skeletal description. Although knowledge representation (KR) systems were not powerful enough when this approach was first proposed by Bobrow and Webber in 1980, the KR system Loom provides the additional needed features: more complete inference of disjointness and inference with respect to coverings. Hovy argued that substantial efficiency improvements over unification-based approaches could be obtained by this classification approach because of inherent structure sharing, consistency maintenance, and explicit type-hierarchy representations. He presented an optimistic view of the future cooperation possibilities between KR and NL researchers, who have traditionally kept their distance. This view is based on the fact that Loom developers have already implemented suggestions useful for NL processing.

J. Hobbs from SRI reported on his continuing search for the optimal point in the efficiency versus power scale for the abduction component of TACITUS. Given that abduction is combinatorially explosive, an empirical investigation is being done exploring 3 different techniques for controlling the computational cost of abduction: the use of a type hierarchy, the unwinding of transitivity axioms, and various heuristics for reducing the branch factor of the search. Large speedups were reported due to the incorporation of a type hierarchy for pre-filtering axioms and eliminating some inconsistent assumptions. Where some transitivity axioms were creating arbitrary recursion, they were "unwound" to limit recursion depth. Finally, several heuristics, such as proving easy (specific) conjuncts first and propagating those results, were presented for the purpose of reducing branching of the search.