

Appeared in the ACL's Finite String newsletter, summer 1986,
issue 12(2), pp. 139-141

The Counselor Project at the University of Massachusetts

David D. McDonald & James D. Pustejovsky
Department of Computer and Information Science
University of Massachusetts,
Amherst, Massachusetts 01003

Participants in the Counselor Project, Fall 1984 through Summer 1986:

Principal Investigators: Edwina L. Rissland, David D. McDonald, Wendy G. Lehnert

Research Associates: Beverly Woolf, James D. Pustejovsky

Graduate Students: Marie M. Vaughan, Brian Stucky, Penelope Sibun, Seth Rosenberg,
Kelly Murray, Kevin Gallagher, JoAnn M. Brooks, John Brolio, Sabine Bergler,
Kevin D. Ashley, Scott D. Anderson

Introduction

The Counselor Project began in the fall of 1984 with the goal of exploring basic problems in discourse structure and text processing within an integrated interface to a strong expert system. The program that we have developed, Counselor, integrates separately developed components for natural language generation (Mumble see [7], [8], [9]), parsing (Plum [5]), and case-based legal reasoning (Hypo [1], [2]). It adds a newly developed component, Cicero ([10]), positioned between the two text processors and the expert system; Cicero is responsible for managing textual inferences ("reading between the lines") by using common sense models of legal events. Counselor can provide advice to an attorney about how to argue cases involving violations of trade secret law in the computer field. The attorney presents the facts of their case to the system, which may ask questions to elicit other facts that it knows to be relevant. The system then suggests lines of argument that the attorney might use, drawing on its library of litigated cases to find ones with analogous dimensions.

At its present state of development, Counselor can handle simple variations on a single scenario, exemplified by the following dialog:

User: *I represent a client named HackInc, who wants to sue SwipeInc and Leroy Soleil for misappropriating trade secrets in connection with software developed by my client. HackInc markets the software, known as Autotell, a program to automate some of a bank teller's functions, to the banking industry.*

Counselor: *Did Soleil work for HackInc.?*

User: *Yes, he was a key employee on the Autotell project.*

Counselor: *Did he later work for SwipeInc.?*

User: *Yes.*

Counselor: *You can argue that there is an implied agreement arising out of Soleil's employment with HackInc. that he not disclose any trade secret information to which he gained access by virtue of his employment.*

Motivations

Consequential results in natural language research will only come from working with a strong underlying program whose communicative needs will challenge the capabilities of state of the art of language interfaces. As a group, we are not interested in building yet another question answering system: our goal is to understand the structure of discourse. We believe that an effective place to begin is with task specific, mixed initiative dialog where the participants' goals cannot be satisfied by single utterances.

Working with a legal reasoning system like Kevin Ashley and Edwina Rissland's Hypo provides particular challenges to natural language research:

- (1) Legal text is structurally complex. The need to avoid ambiguity leads to deeply embedded clauses and heavy noun phrases.
- (2) As both the user and the system have a thorough knowledge of the law, they communicate vastly more information in conversations about legal arguments than ever appears in their literal utterances.
- (3) Hypo's role as an advisory system creates a natural motivation to communicate through language.
- (4) Legal cases are large, complex objects that can be viewed from many alternative perspectives. The purpose for which a case is being described strongly influences which of its attributes are salient and how that information should be structured as a text.

Component Parts

We began the project with three partially developed components, Hypo, Mumble, and Plum, each designed with independent motivations. An initial tension was whether to convert aspects of these programs that did not seem apt in their new setting, or alternatively to interpose new components between them to smooth out the differences. We concluded that the motivations underlying each component were strong enough that we should not change them just because they were now working together.

Hypo reasons with cases and hypotheticals. Actually litigated legal cases are encoded and indexed by "dimensions", which capture the utility of a case for making a particular kind of argument. When evaluating new cases, Hypo first analyzes them in terms of the dimensions they involve. Relevant cases are then retrieved to guide the reasoning. The system may ask pertinent questions about facts now found to be relevant. When the analysis is complete, Hypo describes the arguments available to the user, and responses and counter responses that may follow.

Mumble, the linguistic component for generation, is responsible for realizing conceptual specifications as grammatical text cohesive with the discourse which proceeds it. Mumble works within a description directed framework. Its input specification is a description of the message the underlying program wants to communicate. This description is executed incrementally, producing an intermediate linguistic representation which defines the text's grammatical relations and imposes constraints on further realization. This surface structure description is concurrently executed, producing the actual text.

Plum is a conceptual analyzer which has been given a well defined schematic structure so that it can be easily extended. It parses by doing prediction and completion over semantic concepts implied by the words rather than over syntactic categories. As in other conceptual analyzers, no explicit surface structure is recovered. Plum's output is the set of completed frames.

Cicero is a new component, a discourse and inference manager between the language components and the expert system. From the understanding side, Cicero must integrate the clause by clause output of the parser into the larger discourse context, recognizing, for example, when noun phrases refer to the same object. In interpreting these small, lexically derived frames, Cicero draws on its own representation of events which bridges the gap between the way such information is expressed in language and the way it is organized for expert legal reasoning. For generation, Cicero is responsible for planning the message that is given to the generator. In particular, it determines what information should be included and what may be omitted as inferable, and it selects pivotal lexical items with appropriate perspective and rhetorical force.

Future Directions

While the accomplishments of the individual components of Counselor are interesting in their own right, the greatest effect of the project has been to provide a workbench for studying the problems of language in an integrated context. Perennial problems in anaphora, lexical semantics, aspect, etc. become more tractable in an integrated system where there is a discourse context and intensional motivation. There are also semantic generalizations between the level at which the text processors operate and the level of the expert system which are more easily captured when parsing and generation can be studied in unison. On a larger scale, an explicit discourse manager, a requisite for more complex dialogs, can only be developed once an integrated system exists.

References

- [1] Ashley, Kevin D. (1986) "Modeling Legal Argument: Reasoning with cases and hypotheticals -- a thesis proposal", Technical Report 10, The Counselor Project, Department of Computer and Information Science, University of Massachusetts at Amherst.
- [2] Ashley, Kevin D. and Edwina L. Rissland (1985) "Toward Modeling Legal Argument", Proceedings of the 2nd International Congress LOGICA, INFORMATICA, DIRITTO, Istituto Per La Documentazione Giuridica, Florence, Italy.
- [3] Brooks, JoAnn M. (1985) "Themis: A Discourse Manager", unpublished Master's thesis, Department of Computer and Information Science, University of Massachusetts at Amherst.
- [4] Gallagher, Kevin (1986) "The Design and Implementation of Cicero", unpublished Master's thesis, Department of Computer and Information Science, University of Massachusetts at Amherst.
- [5] Lehnert, Wendy G. and Seth Rosenberg (1985) "The Plum User's Manual" Technical Report 1, The Counselor Project, Department of Computer and Information Science, University of Massachusetts at Amherst.
- [6] McDonald, David D. (1986) "Natural Language Generation: Complexities and Techniques", to appear in Nirenburg (ed.) Theoretical and Methodological Issues in Machine Translation, Cambridge University Press.
- [7] McDonald, David D. and James Pustejovsky (1985) "Description Directed Natural Language Generation", Proceedings of IJCAI-85, pp. 799-805.
- [8] McDonald, David D. and James Pustejovsky (1985) "TAGs as a Grammatical Formalism for Generation", Proceedings of the 23rd Meeting of the Association for Computational Linguistics, pp. 94-103.

- [9] McDonald, David D. and James Pustejovsky (1985) "SAMSON: A Computational Theory of Prose Style for Natural Language Generation", Proceedings of the 1985 meeting of the European Association for Computational Linguistics.
- [10] Pustejovsky, James (1986) "An Integrated Theory Discourse Analysis", Technical Report 11, The Counselor Project, Department of Computer and Information Science, University of Massachusetts at Amherst.
- [11] Rissland Edwina L., Edward Valcarce, and Kevin Ashley (1984) "Explaining and Arguing with Examples", Proceedings of AAAI-84.
- [12] Vaughan, Marie M, and David D. McDonald (1986) "A Model of Revision in Natural Language Generation", Proceedings of the 24th Meeting of the Association for Computational Linguistics.