NLP Research at the Yale AI Project

Natalie Dehn
Yale University

1. The Yale AI Approach to NLP

The Yale AI Project is an interdisciplinary research effort, made up primarily of computer scientists and psychologists, but also including, over the last several years, researchers trained in linguistics, philosophy, and anthropology. Current graduate students are officially affiliated with either the Computer Science or Psychology departments, but are educated, beyond general departmental requirements as Cognitive Scientists.

Natural language research by the Yale AI Project has, from the start, been concerned with (1) conceptual representation (the meaning underlying natural language texts or utterances) and (2) prior knowledge needed in understanding a text or utterance (or assumed, but not explicitly expressed, in its expression). In the last several years, the former has evolved into concern with how understanding is represented in memory (i.e., how it is integrated into what one already knows and hence can effect further understanding), the latter into how the prior knowledge one applies in understanding or generation is represented in such an integrated memory, and how that knowledge is accessed.

Thus, our understanding of "understanding" has shifted. It is no longer a sufficient test of comprehension that a system answer comprehension questions or retell what it has understood as a paraphrase or translation; it is now expected that a system learn from what it's read (or get bored if there's nothing new). Thus a system reading several stories is expected to make use of what it has learned from earlier ones in understanding later ones, and a system understanding a comprehension question on a story it has just read is expected to make use of it's understanding of the story in the very parsing of the question about it. It is also expected that understanding something new lead to spontaneous remindings of things previously understood (and hence stored) in similar terms.

In short, NLP research at Yale AI has merged into research on memory modelling.

2. Yale AI Personnel

Faculty: Bob Abelson, John Black, Drew McDermott, Wendy Lehnert, Chris Riesbeck, Roger Schank, Elliot Soloway

Graduate Students: Valery Abbott, David Atkinson, Bill Bain, Paul Barth, Larry Birnbaum, Mark Burstein, Gregg Collins, Ernie Davis, Livingston Davies, Natalie Dehn, Mike Dyer, Margot Flowers, Tom Galloway, Abraham Butman, Kris Hammond, Eduard Hovy, Pete Johnson, John Leddo, Stanley Letovsky, Steve Lytinen, Beth Marshburn, Rod McGuire, David Miller, Tony Passera, Richard Plevin, Brian Reiser, Scott Robertson, Steven Salzberg, Colleen Seifert, Michael Wilk

Research Staff: Laurence Danlos, Kate Ehrlich, Jim Galambos, Shoshi Hardt, Shun Ishizaki, David Johnson, Steve Slade

Undergraduates: Larry Hunter

3. Projects

Natural language and memory research at the Yale AI Project occurs, in large part, within projects, or smaller subgroups. While there is a common subpool of theoretical issues underlying these diverse projects, and a free flow of ideas and personnel between them, each project has a somewhat different focus or subset of issues of most immediate concern. Most projects attack their issues both through the attempt to construct working process models (AI programs) and through the collection of confirming or disconfirming data (psychology experiments on human subjects).

Each project revolves around the construction of a program modelling human performance of a particular "task", such as in-depth narrative comprehension [BORIS], story invention [AUTHOR], political argumentation [ABDUL/ILANA], personal-problem advice generation [JUDGE], domestic conversation [MAGPIE], learning (through reading of newspaper stories on a topic) [REGAL, ALFRED], or translation [MT]. The degree to which a "task" is itself a central theoretical concern of a project or primarily a vehicle for exploring underlying issues, such as opportunistic planning, failure-driven memory, explanation, reminding, or reconstruction, varies.

4. References

The following is a (greatly abridged) bibliography of recent publications on or underlying natural language research at the Yale AI Project.

Black, John B; Wilkes-Gibbs, Deanna and Gibbs Jr., Raymond W. (1981). What Writers Need to Know That They Don't Know They Need to Know. Cognitive Science Research Report #8. Yale University Cognitive Science Program.

Dehn, Natalie (1981). Memory in Story Invention. Proceedings of the Third Annual Conference of the Cognitive Science Society.

Dyer, Michael G. (1981). Integration, Unification, Reconstruction, Modification: An Eternal Parsing Braid. Proceedings of the Seventh International Joint Conference on Artificial Intelligence.

Johnson, Peter N. and Robertson, Scott P. (1981). MAGPIE: A Goal-based Model of Conversation. Research Report #206. Yale University Department of Computer Science.

Lehnert, Wendy G.; Black, John B. and Reiser, Brian J. (1981). Summarizing Narratives. Proceedings of the Seventh International Joint Conference on Artificial Intelligence.

Lehnert, Wendy; Dyer, Michael G.; Johnson, Peter N.; Yang, CJ; Harley, Steve (1981). BORIS -- An Experiment in In-Depth

Understanding of Narratives. Research Report #188. Yale University Department of Computer Science.

McGuire, Rod; Birnbaum, Lawrence; Flowers, Margot (1981). Opportunistic Processing in Arguments. Proceedings of the Seventh International Joint Conference on Artificial Intelligence.

Schank, Roger C. (In Press) DYNAMIC MEMORY: A Theory of Learning in Computers and People. Cambridge University Press.

Schank, Roger C. and Birnbaum, Lawrence (1980). Memory, Meaning, and Syntax. Research Report #189. Yale University Department of Computer Science.

Research at Xerox PARC

Ronald M. Kaplan
Xerox Palo Alto Research Center

The Xerox Palo Alto Research Center maintains a small, theoretically oriented natural language research program. The principal investigators are Ronald M. Kaplan (Cognitive and Instructional Sciences Group) and Martin Kay (Computer Science Laboratory). For the last several years, we have focussed on linguistic, psycholinguistic, and computational aspects of syntax and morphology.

Syntax:

We have devised two new formalisms for characterizing the syntactic structures of natural language. Kaplan, in collaboration with Joan Bresnan (Department of Linguistics and Philosophy, MIT) has developed the theory of "Lexical-Functional Grammar" (LFG). This formalism has been designed to serve as a medium for expressing and explaining important generalizations about the syntax of human languages and thus to serve as a major vehicle for linguistic research. Of equal significance, it is a restricted, mathematically tractable notation for which simple, psychologically plausible processing mechanisms can be defined.

Lexical-functional grammar represents the convergence of two separate lines of investigation, the psycholinguistic and computational research of Kaplan, Wanner, Woods, and others, which evolved in the Augmented Transition Network tradition, and Bresnan's work within the framework of transformational grammar. The formal notions of LFG are presented in Kaplan and Bresnan (in press). The collection edited by Bresnan (in press) contains a variety of other papers that apply the LFG formalism to a variety of linguistic and psycholinguistic issues. For example, Ford, Bresnan, and Kaplan (in press) provides an explanation for syntactic closure effects in performance based on an LFG competence grammar and a model of LFG parsing derived from Kaplan's General Syntactic Processor. This parsing model has been implemented at Xerox, and we have begun to use it as a tool for grammar debugging in addition to using it to model human psycholinguistic processing.

The second syntactic formalism is Kay's "Functional Grammar" (Kay, 1979). The underlying mechanisms of this formalism are similar to those of LFG: both theories are based on simple abstract operators for manipulating descriptions of functional relations given in declarative, non-procedural specifications. The differences between the formalisms are somewhat superficial but nonetheless significant: Whereas both constituent and functional structures are first-class formal objects in LFG, Functional Grammar de-emphasizes the role of constituent structure, giving proportionately more prominence to functional relationships. Because of this, Functional Grammar is more closely connected to notions in Halliday's Systemic Grammars than to transformational notions. Both parsing and generation algorithms have been implemented for Functional Grammars, and investigations of various computational trade-offs have been investigated.

Morphology and word recognition:

We have been investigating formalisms for describing how canonical word forms vary when they are composed with other words and affixes in text and speech. Our goal in this area also is to arrive at a framework that permits succinct statements of various morphophonemic and morphographemic processes but which also can be mapped in a systematic way into efficient procedures both for recognition and generation. We have been concerned not just with the relatively impoverished set of word variations of English, for which simple suffix-stripping algorithms are usually used, but with phenomena that occur across the languages of the world (e.g., vowel harmony, umlauting). We have found that if the individual rules in a standard phonological rule formalism are restricted so that the non-contextual part never applies to its own output, then they can be modeled by symmetric finite-state transducers. Furthermore, if the rules of a morphophonemic or morphographemic grammar are applied non-cyclicly, then the set of individual transducers can be composed into a single one that models the whole grammar (Kaplan & Kay, 1981). This transducer can be combined with certain dictionary lookup algorithms to yield procedures that are far more economical for both recognition and production than any strategies using ordered rules directly (Kay & Kaplan, forthcoming).

References:

Bresnan, J. (ed.) The mental representation of grammatical relations. Cambridge: MIT Press, in press.

Ford, M., Bresnan, J., & Kaplan, R. M. A competence based theory of syntactic closure. In Bresnan (in press).

Kaplan, R. M. & Bresnan, J. Lexical-functional grammar: A formal system of grammatical representation. In Bresnan (in press).

Kaplan, R. M. & Kay, M. Phonological rules as finite state transducers.