Final Report on a Visit to ICOT

Carl J. Pollard

Center for the Study of Language and Information
Stanford University
Stanford, California, U.S.A

and

Hewlett-Packard Laboratories (Consultant)
Palo Alto, California, U.S.A.

1. Introduction

From July 21, 1986 to August 1, 1986, I was invited to ICOT as a visiting researcher in the field of natural language processing (NLP). Given the relatively brief duration of the visit, no specific research objectives were set; rather, the main purpose was to provide a forum for information exchange and discussion concerning our current research on unification-based approaches to natural language syntax and semantics. On the one hand, I was brought up-to-date on many of the NLP-related projects under way at ICOT through a series of demonstrations and briefings by members of ICOT's Second Research Laboratory; on the other hand, I was provided with opportunities to familiarize ICOT researchers and other interested parties with my recent work on the grammatical theory Head-Driven Phrase Structure (HPSG) and related issues concerning the logical foundations of linguistic theory, through a series of tutorial lectures. Less formally but of no less importance, there was ample opportunity in between the scheduled presentations and over lunch for the kind of free-form brainstorming that is an indispensable ingredient of true collaborative research.

2. Demonstrations and Briefings by ICOT Researchers

I was able to get some grasp of the remarkable scope of ICOT's research mission through a series of presentations by members of ICOT and the associated working groups. In addition to a helpful introductory briefing on the overall goals and structure of the Fifth Generation Computer Systems (FGCS) project and a tour of the facilities, I received more detailed presentations on ongoing NLP-related work, including: the Japanese Phrase Structure Grammar (JPSG) grammar-parser system; the CIL language for programming with complex indeterminates; applications of CIL to discourse understanding (the DUALS system) and the interpretation of honorifics; parallel parsing; and the use of Japanese as a program specification language. Since I am trained primarily as a linguist, secondarily as a mathematician, and scarcely at all as a computer scientist, this wide array of presentations turned out to serve a correspondingly wide variety of purposes: in some cases, it helped bring me up-to-date on

current work very closely related to my own, to perceive similarities and parallels as well as divergences in both general approach and matters of technical detail; in other cases the function was more to give me a bird's-eye view of a line of work outside my area of expertise. These presentations contributed greatly to my understanding of what ICOT's work is really about.

3. Presentations to ICOT and Associated Working Groups.

Much of my time at ICOT was spent preparing and delivering some 12 hours of presentations on my own linguistic research. My first presentation was in the form of a lecture to the Natural Language Processing System Working-Group on some current developments in unification-based approaches to natural language. The emphasis in that lecture was upon the logical foundations of linguistic theory, with a brief sketch of the overall structure of the theory HPSG.

The remainder of my presentations were in the form of a three-session sequence of tutorials in which I gave a more detailed introduction to HPSG, with the emphasis on particular linguistic phenomena. The content of these sessions were as follows:

First Session: methodological and philosophical preliminaries concerning the information-processing perspective on language and the unification-based approach; kinds of linguistic information; the structure of grammatical categories and lexical entries;

Second Session: rules of grammar and principles of universal grammar, and the treatment of particular phenomena such as subcategorization, government, complementation, and inversion;

Third Session: more detailed treatment of particular phenomena, such as control, agreement, and unbounded dependencies, as well as comparison with corresponding accounts in other linguistic theories (Transformational Grammar, Lexical-Functional Grammar, and Generalized Phrase Structure Grammar).

The tutorials were attended by 30-40 people, including ICOT members as well as interested parties from other organizations in and around Tokyo, and allowed ample time for (often lively) discussion.

4. Informal Discussion

As valuable as the formal presentations described above have been for me (and, I hope, to my hosts), I am grateful to the organizers of my visit for leaving plenty of time for unstructured discussion, for this is the kind of interaction during which I think most new ideas really get generated. During my visit, such interaction took the form of free group discussions and more intensive one-on-one interactions with individual researchers. During the first few days of my visit, several sessions were scheduled for free discussion with the JPSG Working Group. The concerns of this group -- the encoding and parsing of natural language grammar using a unification-based phrase-structure formalism -- is directly related to my own research interests, and I think these extended sessions provided an excellent opportunity for both sides to compare and contrast our research from the point of view of goals, methodology, and overall approach. After the orientation

provided by such discussion (and the briefings and demonstrations mentioned above), it was then easy to settle down into ongoing discussions with individual researchers on technical topics, e.g. plans for treating a wider range of grammatical rule types, prospects for opening up the relatively unmined field of morphology, differing strategies for augmenting pure unification with schemes for maintaining functional relationships between still-uninstantiated variables, comparison of representational ("neo-Fregean") versus nonrepresentational (Barwise-Perry) accounts of indexicality, etc.

5. Concluding Remarks

My most lasting impression of the work on NLP going on at ICOT will be my surprise at discovering the remarkable degree to which it has kept abreast of the very latest theoretical developments with which I am acquainted. Indeed, given the vague image in my mind before I came to Japan of ICOT as a fundamentally economically-oriented agency charged with keeping Japan in the forefront of technological development, it came as something of a shock to me to realize the extent to which individual members of ICOT and the associated working groups are concerned with issues of theoretical interest; that is, ICOT is far more basic-research oriented than I (and probably the majority of American researchers) imagine it to be. My surprise at this ability to stay current with recent theoretical advances is best exemplified by my discovery that the new book on situation theory by Barwise and Etchemendy, of which I had received a copy only the day before I came to Japan, appeared to be ubiquitous in the Second Research Laboratory, and indeed was the subject of an ongoing and well-attended seminar ten steps away from my desk! Similarly, I found a degree of familiarity with ongoing research in syntax and semantics that allowed me to interact with ICOT workers at the same level of sophisitcation that I am accustomed to at CSLI or Hewlett-Packard. The people I met here are working on many of the same problems, and sharing much of the same background, as my colleagues at home.

Since I am not a computer scientist, I am in no position to speculate about ICOT's prospects for success in the extremely ambitious FGCS project. Within the more limited arena of NLP, my impression is that ICOT's efforts in this area are in a class with the best work that I am familiar with at home; still, it is difficult for me to go away from ICOT without wondering whether it has set its sights too high. According to the master plan, if I understand the big picture correctly, all the bits and pieces of research that I have seen in the past two weeks are destined to be assembled someday into a single integrated technology that will represent a quantum leap beyond anything that now exists. If so, then success depends upon the mutual consistency of all the pieces, and, moreover, upon the success of the new-generation parallel hardware, inference engines, and knowledge representation technology; but no one knows yet the extent to which these underlying technologies will succeed.

Moreover, it is not clear to me that the individual pieces of work that I have seen actually will lend themselves to being easily merged together: to take a very simple example, the basic mechanisms for augmenting unification that underly the JPSG parser on the one hand and the CIL freeze predicate on the other hand appear to be based on mutually inconsistent assumptions, and I cannot help wondering how difficult it will be to hammer out a compromise that will be

satisfactory to all parties. Another related concern is the following: many of the small pieces that I have seen work well in isolation: e.g. the JPSG parser deals well with a fragment of Japanese syntax involving verb complementation and the reflexive pronoun, and some interesting discourse-related problems have been attacked within CIL; but the current JPSG implementation deals with semantics at a trivial level, with discourse not at all, while the discourse-related work was based upon a fairly trivial syntax; moreover, entire broad areas seem not to have been tackled at all yet, e.g unbounded dependencies, scope quantifiers, general anaphora, ellipsis, intensionality, comparatives, etc. The question is whether all these things can be forged together and still run in real time.

Notwithstanding this uncertainty about the ultimate success of FGCS, I think the individual strands of work going on at ICOT are important in their own right, and I hope that something like ICOT will continue to be supported in Japan after FGCS comes to an end (whether or not it achieves its stated goals). Work of this caliber is possible only with the dedication of substantial funds on an ongoing basis, and with institutional cooperation across the government, academic, and industrial sectors. The degree of organization and support required, therefore, is too great for it all to rise or fall with the success of one project. Experience seems to teach that the benefits of research cannot be predicted; science goes forward, but not in directions that anyone can predict. To conclude on a note that is both optimistic and cautionary: ICOT's work may well fall short on some fronts of the ambitious goals that have been set for it, but like world-class scientific research everywhere, its ultimate impact in unforeseen areas will be considerable.

6. Acknowledgements

So many people have contributed in one way or another to making my brief stay at ICOT both profitable and enjoyable that it is difficult to give everyone involved with my visit due credit. Among those whose names come to mind, I would first like to thank ICOT's director, Dr. K. Fuchi, and T. Yokoi, Chief of the Second Research Laboratory, for honoring me with their invitation. I owe special debts of gratitude to Mr. H. Kusama for his extraordinary kindness and thoroughness in arranging my visit, and to Mr. H. Miyoshi for his unflagging patience and attention to detail in seeing that all my day-to-day on-the-job needs were taken care of. I am grateful, too, to Prof. H. Tanaka for providing me with the opportunity to present some of my ideas to the Natural Language Processing System Working-Group. For providing me with ongoing and intense intellectual stimulation, I am particularly indebted to Prof. T. Gunji and Dr. K. Mukai; of less importance to ICOT, perhaps, but certainly of no less importance to me, I would like to thank them for their hospitality in Osaka and Kamakura respectively, and for showing me that there is more in Japan than a research institute and ten subway lines! Among the other researchers, thanks are due to Mr. Y. Matsumoto, R. Sugimura, and Y. Tanaka -- as well as several others whose names I cannot recall -- for taking the trouble to present me with briefings and demonstrations on ongoing research. Finally, I would like to thank some of the people who provided me with intellectual and personal companionship, both on the job and over beer, especially Y. Harada, K. Hasida, and H. Sirai. All of these people helped to make my visit intellectually rewarding and to lend it a real human texture as well.

RESUME OF LINGUISTIC RESEARCH

Carl J. Pollard did graduate work in mathematics (M.S. Brown University, 1973) and linguistics (Cornell University; M.A. Indiana University, 1978; Ph.D. Stanford University, 1984). His graduate work in linguistics at Cornell and Indiana concentrated upon Chinese historical phonology and dialectology; later at Stanford, working under the direction of Ivan Sag and Thomas Wasow, his focus of interest shifted to syntactic and semantic theory and their mathematical foundations.

In his doctoral dissertation, entitled Generalized Phrase Structure Grammars, Head Grammars, and Natural Language, he defined a class of generalizations of context-free grammars (head grammars), described some of their mathematical and computational properties, and used them as the formal basis for a lexically based and computationally tractable theory of grammatical structure. Since the completion of his doctoral studies, he has taught in the Linguistics Department of Stanford University, undertaken research as an associate of Stanford's Center for the Study of Language and Information (CSLI), and served as a linguistics consultant at Hewlett-Packard Laboratories.

At CSLI, Pollard's research has centered upon the development of Head-Driven Phrase Structure Grammar (HPSG), an information-based theory of natural language syntactic and semantic structure capable of integrating and synthesizing insights and results produced by a variety of current syntactic-semantic approaches, including Categorial Grammar, Generalized Phrase Structure Grammar, Lexical-Functional Grammar, Situation Semantics, and Discourse Representation Theory. In common with computational linguistic formalisms such as Martin Kay's Functional Unification Grammar and SRI International's PATR-II, HPSG describes (represents, classifies) linguistic structures declaratively, in terms of an information domain consisting of directed graphs or attribute-value matrices (AVM's) whose principle combinatory mode is the recursive information-merging operation called UNIFICATION.

Recent HPSG research has focused upon foundational issues, including: (1) the mathematical, computational, and semantic properties of the AVM formalism in terms of which the theory is couched; (2) the precise theoretical status of lexical entries, grammar rules of particular languages, and principles of universal grammar; and (3) the prospects for bringing HPSG theory within the compass of current efforts by a research group at CSLI to provide axiomatic foundations for situation theory.

In addition to purely theoretical matters, Pollard's work has also been concerned with the implementation of natural language computer systems based upon sound linguistic principles. For the past several years he has consulted with Hewlett-Packard's Natural Language Project on the development of an HPSG-based linguistic system. He is also involved with a current drive by researchers at Stanford and SRI International to construct a new computational linguistics development environment. Written in Common LISP, and using SRI's PATR-II as a springboard, this new effort seeks to produce a pedagogical and research tool that takes advantage of recent theoretical results in both linguistics and computer science.