Directed Learning in a Complex Domain. Research Report (forthcoming) Department of Computer Science, Yale University

[18] Schank, R S, Lehnert, W. (1979) The Conceptual Content of Conversation Research Report #160 Department of Computer Science, Yale University.

[19] DeJong, G F (1979) Skimming Stories in Real Time: An Experiment in Integrated Understanding Ph.D. Thesis Department of Computer Science, Yale University

[20] Abelson, R P (1980) Differences Between Belief and Knowledge Systems Research Report #1 Cognitive Science Program, Yale University

[21] Abelson, R P (1980) The Psychological Status of the Script Concept Research Report #2. Cognitive Science Program, Yale University

[22] Black, J B, Bern, H (1980) Causal Coherence and Memory for Events in Narratives. Research Report #3 Cognitive Science Program, Yale University

[23] Black, J B (1980) Psycholinguistic Processes in Writing Research Report #4 Cognitive Science Program, Yale University

[24] Abbott, V, Black, J B (1980) The Representation of Scripts in Memory Research Report #5 Cognitive Science Program, Yale University

[25] Abelson, R P (1980) In Defense of Common Sense Representations of Knowledge Research Report Cognitive Science Program, Yale University

---

AI at AI&DS

edited by Brian P McCune

Company Overview

Advanced Information & Decision Systems (AI&DS) is a relatively new, employee-owned company that does basic and applied research, product development, and consulting in the fields of artificial intelligence, computer science, decision analysis, operations research, control theory, estimation theory, and signal processing. AI&DS performs studies, analyses, system design and evaluation, and software development for a variety of industrial clients and government agencies, including the Departments of Defense and Energy. The AI&DS technical staff has training and experience in mathematical sciences and engineering, with most members holding advanced degrees. The staff is augmented by a number of university experts who serve as consultants.

AI&DS is located in the San Francisco Bay area near Stanford University. An in-house computing facility for symbolic, numeric, and word processing, based on a DEC VAX-11/750, is on order. Current computing resources include numerous hardcopy and display terminals and a Cromemco System Three with Digital Graphic Systems CAT-200 for displaying grayscale images and local printing. Major computation is done via remote access to a number of DECsystem-10s and DECsystem-20s on the ARPAnet, as well as a VAX-11/780.

AI Interests

Two major themes of AI&DS are relevant to its AI interests. One theme is *technology transfer*. AI&DS has assumed the role of an agent for transferring technology from the basic research found primarily in universities to the applications found in government and industry In this role AI&DS is concentrating on applied research, while maintaining strong components of basic research and development as well

The second theme is *interdisciplinary research*. For example, AI, control theory, and decision analysis are three disciplines that often attack similar problems, but with different tools and under different assumptions about the nature of the problems. As examples, all three disciplines involve aspects of hypothesis formation and decision making, and all three may use search trees as a tool AI&DS is involved in projects (e.g., the last three described in the next section) to understand when one technology is better than another and how technologies may be fruitfully combined.

AI is currently applied in seven AI&DS research projects in five overlapping areas: software aids, image understanding, hypothesis formation, distributed AI, and AI for decision making. In addition, AI&DS has an interest in the areas of knowledge-based system development for various applications, advanced user interfaces, robotics, integration of information from multiple sources [Drazovich & Wishner-81], and fuzzy aids for control [Tong-80] and decision making [Tong & Bonissone-80]

Current AI Research Projects

This section briefly describes each research project at AI&DS that has some AI component. After the title of each research project, the principal investigator is listed.

Advanced Tools for Software Maintenance
Brian P McCune

This research is to evaluate the need for and to design prototype tools that apply advanced software technologies to the problem of maintaining large software systems written in ADA. The technologies being considered include artificial intelligence, automatic programming, program verification, very high-level languages, program transformations, program optimization, and interactive programming aids.

The effort is directed primarily toward tools that will improve productivity, increase reliability, and lower costs during the maintenance phase of the software life-cycle. One such tool is an intelligent program editor that understands the syntax and semantics of ADA and something of the application domain as well. Another tool is a maintenance manager, an expert system for keeping track of what needs to be done next (e.g , what tools to invoke, what files to update, what permissions to receive). In order for the most appropriate tools to be proposed and evaluated, the study also involves an analysis of current software maintenance practices and problems [McCune-80B].

Program Reference Language
Brian P McCune

The general scientific goal of this project is to improve our fundamental capability for manipulating computer programs We want to understand the theory and techniques for intelligently navigating (moving through) a database of programs, manipulating programs, and automatically monitoring the process for errors. This understanding will allow us to develop intelligent aids for computer program acquisition, especially an advanced program editor.

Specifically, the project is to design a program reference language for navigating through a database of programs [McCune-79]. The pattern matching primitives of this language provide a simple, uniform framework for accessing a piece of a program by a combination of textual, syntactic, contextual, historical, semantic, and pragmatic indices. In order to test the utility of the proposed constructs of the program reference language, the program reference language specifications will be embedded within a set of commands in an intelligent program editor for languages such as PASCAL and ADA When completed, this editor will aid a programmer by (1) allowing program parts to be referred to by the use of techniques specific to programs, (2) allowing the editing or manipulation of program parts by the use of techniques specific to programs, and (3) incrementally checking for errors of semantic completeness, consistency, and ambiguity [McCune-80A]

By automating many mundane programming chores and by catching programming errors earlier than ever before possible, the results of this research will impact the software problems of reliability, cost, production time, programmer productivity, frequency of change, increasing complexity, and ease of programming by nonexperts Scientific contributions will be made at the intersection of the disciplines of programming languages, software engineering, artificial intelligence, and database management. The specific technical areas being addressed include meta-programming language design, program semantics, programming environments, program editors, programming methodology, automatic programming, interactive artificial intelligence, pattern matching, knowledge representation, and information retrieval.

Radar Image Classification
Robert J Drazovich

This investigation is applying visual image understanding (IU) and other AI techniques to the problem of classifying ships from high resolution radar imagery. During the feasibility analysis, a small database of high-level ship information was created, key feature extraction and classification techniques from recent IU work were modified to work in the domain of radar images (which are more distorted and have more highly variable feature reflectance than visual images), alternate hypothesis representations and system organizations were considered, and the feasibility of a classification system was demonstrated by trying the techniques on a few images using a bottom-up control structure [Drazovich, Lanzinger, & Binford-81A; Drazovich, Lanzinger, & Binford-81B]

Design and construction of a complete prototype classification system are now under way. A multi-level hypothesis structure will be used [Drazovich, Brooks, & Foster-79], with the levels representing (from lowest to highest) radar signals, images, image features, surface features, ship features, and ship identity. Low-level extraction routines will first create a symbolic representation of the ship in terms of image features. Then classification rules, in part extracted from expert radar image analysts, will be used for the final classification. These rules relate the image features to important ship features (e.g., length, width, superstructure placement, decks). The control structure will allow for top-down, predictive processing, as well as bottom-up The sensitivity of the system to using contextual information from sources other than the current image will also be studied.

Characteristics of Hypothesis Formation Systems
Robert J Drazovich

The long-term goal of this project is to create a methodology by which (1) the feasibility of a knowledge-based sytem for solving a new hypothesis formation problem can be assessed and (2) if deemed feasible, appropriate forms of hypothesis representation, knowledge base, and reasoning process may be recommended. This (non-automated) approach is to be contrasted with that of domain-independent hypothesis formation systems that are restricted to using one or a small number of techniques.

The current goal is to characterize the space of successful knowledge-based hypothesis formation systems To this end, the effort has developed a simple definition of such systems in terms of input data, output (hypothesis), and processing (reasoning and supporting knowledge base), a set of appropriate systems for characterization; a detailed, though still tentative, outline of the important characteristics of hypothesis formation systems; preliminary characterizations of the SIAP, MYCIN, DENDRAL, and CENTAUR systems; and preliminary observations based upon the similarities and differences of the four systems [Drazovich, McCune, & Buchanan-81]

Distributed Hypothesis Formation for Distributed Sensor Networks
Richard Wishner

In this effort we are developing and evaluating alternate distributed hypothesis representations, formation techniques, and evaluation techniques for distributed networks of sensors. By distributed techniques we mean, for example, schemes analogous to the ARPAnet routing protocols, wherein a node receives information only from its neighbors and computes the correct first segment of a multi-segment path that is good (or even optimal under some criterion) when viewed at the level of the whole network From a global perspective each node has incomplete information, so an assessment of the total situation requires information from all nodes. However, each node, by using only local information from its neighbors, assesses enough of the situation to do the right thing.

In our problem, each node of the network has a processor, some local sensors (e g , acoustic, radar), and limited communications with its neighbors Sensors at different nodes may have overlapping coverage The goal is to develop techniques so that, in a distributed manner, all nodes can obtain and maintain a global hypothesis about the situation throughout the geographic area covered by the network A key research issue is that of how one node formulates hypotheses about the situation throughout the geographic area covered by the network A key research issue is that of how one node formulates hypotheses about the situation using as inputs only information from its own sensors and summarized hypotheses provided by neighboring nodes. This can be viewed as a knowledge-based hypothesis formation problem in which information of multiple types and from multiple sources must be integrated over time. Another important issue is that of how information can best be propagated through the network, including determining what information should be transmitted and what the best control structure is (e.g., to avoid information looping) Some possible techniques have already been applied to a non-distributed problem [Wishner et al.-81].

Distributed Decision Making Environments
Richard P Wishner

This project is to develop a methodology for mapping distributed decision making problems into an appropriate environment, in terms of decision making techniques and their implementation We are characterizing the types of problems that are amenable to a semiautomated, distributed decision making approach, and we will create detailed scenarios to exemplify a few such problems. Candidate decision making techniques include those from distributed AI, distributed control, decision analysis, and decision support systems. Implementation issues include the appropriate network architecture, node architectures, and the distribution of data Another important issue is the appropriate roles of humans and machines in the decision making process.

Option Generation Techniques for Decision Making in Crises
J. Roland Payne

This effort is concerned with the problem of generating options for courses of action in unique decision problems. Unique problems are those with no direct historical precedents (e.g., the Cuban Missile Crisis, the taking of the U.S. embassy in Iran) and for which, therefore, no predeveloped plans can be implemented. Such problems are characterized by a lack of understanding of all the problem components, and insufficient time to permit thorough and exhaustive information collection and analysis. The criticality of the problem produces a compounding negative effect on the decision makers because of cognitive stress and degraded judgment.

Our approach to option generation starts with an initial structuring of the problem in as specific terms as possible This is done by instantiating one or more frames from a library of generic crises created by an analysis of historical crises This library will be organized as a tree of frames, from a root frame representing all crises to more and more specific classes and subclasses of crises further down the tree

Once the initial structuring has been done interactively with the user, a systematic, hierarchical priority assessment is made of all decision factors [Arbel-81]. The approach is structured in a way that accommodates the effect of cognitive biases and cognitive styles on the quality of the decision making process. This approach helps the decision maker to focus quickly on the most important factors, use his creativity to identify other important factors, and iteratively reevaluate the crisis and generate new options

For more information about AI&DS or any of the above research projects, contact:

Brian P McCune
Advanced Information & Decision Systems
201 San Antonio Circle, Ste 286
Mountain View, Calif. 94040

Phone: 415/941-3912

ARPAnet message: MCCUNE@USC-ISIC

References

Arbel-81 Ami Arbel, *Hierarchical Value Assessment in a Task Force Decision Environment*, technical report, Advanced Information & Decision Systems, Mountain View, Calif , January 1981

Drazovich & Wishner-81 Robert J Drazovich and Richard P Wishner, "Artificial Intelligence Approaches to Information Fusion," proceedings of the Air Force Communications and Electronics Association Intelligence Symposium, Rome Air Development Center, Griffiss Air Force Bse, New York, May 1981

Drazovich, Brooks, & Foster-79 Robert J Drazovich, Scottie Brooks, and Scott Foster, "Knowledge-Based Ship Classification," *Applications of Image Understanding and Spatial Processing to Radar Signals for Automatic Ship Classification: Proceedings of a Workshop*, Naval Command, Washington, D.C , February 1979, pp 100-105

Drazovich, Lanzinger, & Binford-81A Robert J Drazovich, Franz X Lanzinger, and Thomas O Binford, *The Feasibility of Artificial Intelligence Techniques for Radar Target Classification*, technical report, Advanced Information & Decision Systems, Mountain View, Calif., January 1981.

Drazovich, Lanzinger, & Binford-81B Robert J Drazovich, Franz X Lanzinger, and Thomas O Binford, "Radar Target Classification," proceedings of the IEEE Pattern Recognition and Image Processing Conference, Dallas, Texas, August 1981, to appear

Drazovich, McCune, & Buchanan-81 Robert J Drazovich, Brian P. McCune, and Bruce G Buchanan, *Characteristics of Hypothesis Formation Systems*, technical report, Advanced Information & Decision Systems, Mountain View, Calif , May 1981

McCune-79 Brian P McCune, *Building Program Models Incrementally from Informal Descriptions*, Ph D thesis, Memo AIM-333, Report STAN-CS-79-772, Artificial Intelligence Laboratory, Computer Science Department, Stanford University, Stanford, Calif , October 1979

McCune-80A Brian P McCune, "Incremental, Informal Program Acquisition," *Proceedings of the First Annual National Conference on Artificial Intelligence*, Stanford University, Stanford, Calif , August 1980, pp 71-73

McCune-80B Brian P McCune, *Analysis of Air Force Software Maintenance*, technical report, Software Technology Program, Advanced Information & Decision Systems, Mountain View, Calif , December 1980

Tong-80 Richard M Tong, "Some Properties of Fuzzy Feedback Systems," *IEEE Transactions on Systems, Man and Cybernetics*, Vol SMC-10, No 6, June 1980, pp 327-330

Tong & Bonissone-80 Richard M Tong and P P Bonissone, "A Linguistic Approach to Decision Making with Fuzzy Sets," *IEEE Transactions of Systems, Man, and Cybernetics*, Vol. SMC-10, No 11, November 1980, pp 716-723.

Wishner et al.-81 R P Wishner, F Lanzinger, H Mesiwala, and S Mori, *Advanced Techniques for Multi-Target Track Formation*, technical report, Advanced Information & Decision Systems, Mountain View, Calif , April 1981.

International Joint Conference on Artificial Intelligence Publications

Proceedings, Sixth International Joint Conference on Artificial Intelligence. Tokyo, Japan. 20-24 August 1979. 2 volumes, 1146 pages. Copies available for $25, payment in advance, book rate; add: $5 for 1st class—US, Canada, Mexico; $20 air parcel post elsewhere. California residents add sales tax Make checks payable to IJCAI-79. Order from:

IJCAI-79
Computer Science Department
Stanford University
Stanford, CA 94305 USA

Proceedings, Fifth International Joint Conference on Artificial Intelligence. Cambridge, Mass., USA, 22-25 August 1977 2 volumes, 1005 pages Copies available for $25, payment in advance, book rate; add: $4 for 1st class—US, Canada, Mexico; $20 air parcel post elsewhere Make checks payable to IJCAI-77. Order from:

IJCAI-77
Department of Computer Science
Carnegie-Mellon University
Pittsburgh, PA 15213 USA

Advance Papers, Fourth International Joint Conference on Artificial Intelligence, Tbilisi, Georgia, USSR, 3-8 September 1975 2 volumes, 944 pages Full size copies available: Volume 1, $129 50; Volume 2, $100 10 Order number 2010729. Order from:

University Microfilms International
PO Box 1467
Ann Arbor, MI USA *or*

University Microfilms International Ltd
18 Bedford Row
London, WC1R 4EJ, England UK

Advance Papers, Second International Joint Conference on Artificial Intelligence Stanford, California, USA, 20-23 August 1973 1 volume, 703 pages Copies available for $15, payment in advance, book rate Make checks payable to SRI International. Order from:

Artificial Intelligence Center, J2044
SRI International
Menlo Park, CA 94025 USA

Advance Papers, Second International Joint Conference on Artificial Intelligence, London, England, UK, 1-3 September 1971 1 volume, 697 pages Full size copies available for $145; microfilm copies, 35 mm positive, $72 50 Order number OP 65449. Order from:

University Microfilms International
(Addresses above)

Proceedings of the International Joint Conference on Artificial Intelligence Washington, DC, USA 7-9 May 1969 1 volume, 715 pages Full size copies available for $145. Order number 2004394 Order from:

University Microfilms International
(Addresses above)