Enclosure 1

RAND'S PROGRAM IN ARTIFICIAL INTELLIGENCE IN MEDICINE

Rick Hayes-Roth
Program Director, Information Processing Systems
Rand Corporation

October 1977

SUMMARY:

This proposal for pilot status in the AIM project includes four sections. The first discusses our idea for applying AI to health services research and a long-term significant commitment to this area. The second section discusses some specific issues in information processing psychology. Rand's commitment to and experience with research community relations is discussed in the third section. The final section briefly discusses other research in whose context the medically oriented research exists. A partial bibliography of relevant work is attached as an appendix.

I. AI IN HEALTH SERVICES

Initially, we are interested in investigating the possibility of applying AI techniques to the problem of health services delivery and related research. Specifically, as an adjunct to Rand's 8-year HEW experiment on National Health Insurance alternatives, we would like to develop a heuristic system to adjudicate claims for reimbursement. There are two qualitative policy problems, coverage and quality of care. The coverage problem is to decide, given a computable representation of the administrative policies underlying a health insurance plan, whether a particular treatment is covered by insurance. For example, would jacuzzi treatments be covered under "physical therapy" or would prophylactics be covered under "family planning." The quality of care problem concerns assessing the validity of particular treatments in the context of supposed illnesses. For example, administration of penicillin for colds is disallowed while purchase of aspirin is reasonable.

In overview, Rand's proposed approach to this problem combines a number of AI methods developed in other projects. We propose combining the rule-directed heuristic reasoning structure from MYCIN with a variety of natural language processing techniques. Knowledge necessary to solve the problem is of three types. First, there is an explicit body of expert rules that are written in official administrative policy handbooks. Second, there is an explicit body of knowledge about language and medicine, that relates the simple natural language expressions used in claims to its semantic meaning and the somatic conditions of the patient. Third, there is a body of implicit, heuristic knowledge that relates special cases to general ones and exceptions to general rules. For example, "if x was purchased without a prescription but the doctor usually provides a prescription for x when it is medically indicated, then with confidance (.6) x was not indicated by the physician and should not be reimbursed." Our goal is to choose an initially manageable subset of knowledge and suitably constrain the task complexity to ascertain whether heuristic methods for claims decisionmaking is feasible. If the initial experiments are positive, we will attempt to undertake a major program to realize such a system in time to accompany the health insurance study final report (1983).

To represent knowledge, we adopt the following conventions. Policy rules are presented in a restricted form of English syntax and these are converted into production rules. To expand the power of production system reasoning capabilities, we represent both rules and category descriptions in specialization-generalization hierarchies. Rules are attached to the hierarchy nodes of maximum generality. Each claim is analyzed as possessing a number of characteristics, and these are associated with specific nodes in the category hierarchy. The most plausible, highest-rated rules that are applicable to the specific nodes or their more general categories are applied in priority ordering to the case characteristics. Decisions are reached by finding the most plausible decision chains that cover the characteristics presented. As cases are adjudicated, explanations of the decisions are produced in a restricted English format.

Originally we will encode claims in a restricted English format to minimize the difficulty of parsing them and categorizing their characteristics. Only when we are confident of the basic strategy will we relax our language processing restrictions. While the ungrammaticality and errorfulness of natural claims language is problematic, we have already developed several methods for handling a large number of the most typical kinds of problem. (That work was developed in the context of the Hearsay-II speech understanding system.)

The benefits which would be realizable by a successful research endeavor are suggested below:

The legislation or rules underlying the benefits coverage would be completely operationalized and visible.

Administrative policies would be operationalized and visible.

Application of legislation and policy would be consistent.

The size and power of the administrative bureaucracy would be constrained.

Costs would be reduced.

The system could be accessed for additional functions, such as deciding if some medical operation under consideration would or would not qualify for reimbursement.

Basic statistics on type and cost of treatments would be readily compilable.

Decentralized computer systems would permit decentralized management, while the common decision programs they execute would insure centralized policy formulation.

Our plans for participation with SUMEX-AIM on this project are as follows. First, we may wish to conduct preliminary computing experiments using INTERLISP and, perhaps, EMYCIN. Some time in early FY78, however, we plan on preparing a proposal for NIH or HEW to support a major research effort along the proposed lines. At that time, we may request sufficient support to purchase a more accessible segment of computing cycles to permit rapid development during ordinary business hours. However, we expect our proposal to be strengthened by the back-up availability of SUMEX-AIM computing resources.

More significantly, however, we are now considering the possibility that Rand could become another node on the SUMEX network or otherwise promote its interest in developing a significant level of effort aimed at the application of AI to health services. Reimbursement claims is simply the tip of a large and important iceberg of difficult issues in the delivery of quality medical care. Since Rand's mission is to improve the functioning of government and since the delivery of health services is an ever-increasing and significant component of the federal budget, it would seem highly desirable to catalyze a major effort to apply AI techniques to its difficult issues. These include determination of quality treatment, utilization review, reimbursement, medical training, patient education, privacy and security, research and evaluation. These areas are largely neglected by current AI researchers, but the potential for a significant symbiotic relationship between advanced computer science and health service analysts is clear. We are interested in assessing the magnitude of the potential impact and moving boldly in this direction if the indications are positive. Working with the present AIM community should provide a valuable source of insight into the workings of a national resource used cooperatively for common programmatic goals.

II. INFORMATION PROCESSING PSYCHOLOGY

Rand is currently concerned with a limited number of issues in cognitive psychology that may be related to mental retardation and speech pathology. I will discuss some of these research topics and recent history in brief.

One area of interest concerns the level of abstractness of memory representations of meaning. While many conventional theorists have postulated (almost taken for granted) the need for a semantically abstract representation of knowledge, our research to date with both sentences and entire texts suggest that meaning-based operations are mediated by superficial codes of remembered information. Originally, Hayes-Roth & Hayes-Roth tested the notion that identical conceptual elements would interfere in memory much the way that identical words interfere with learning and recall of different sentences in which they appear. If memory for meaning were conceptually abstract, such interference would be expected identically in both cases. However, no interference was observed between conceptually related materials. Since then, Hayes-Roth & Thorndyke have further investigated the prominence of lexical information in memory representations of meaning. They have found that manipulations as trivial as replacing common words in different sentences both eliminates interference and, at the same time, makes virtually impossible the integration of diverse sources of information for syllogistic reasoning.

Because of the difficulty of accounting for semantically based behavior without the introduction of assumed, intermediate, abstract conceptual entities, it is important to develop an alternative theory of conceptual information processing. Following Quillian's seminal work, Hayes-Roth & McDonald have recently developed a semantic language processor that employs only lexical tokens and relations. The knowledge necessary for language understanding is taken directly from a dictionary, and these definitions are converted to a semantic net by a simple production system. Interpretation of phrases and clauses is performed by spreading excitation semantic net intersection searches that emanate in parallel from the nodes corresponding to words in the phrase. Syntactic constraints are realized through limitations placed on this intersection search that are consistent with the sequential ordering of phrase constituents.

Obviating intermediate theoretical constructs such as semantic primitives is an important objective, if it can be realized. Understanding how the lexically based semantic network can parsimoniously explain a variety of behavioral data is the essential step in the evolution of a scientific theory of the role of language in thinking and an understanding of why perceptual characteristics of stimuli can dominate and preclude their appropriate conceptual processing. Jim Miller's proposal should be interpreted within this larger, theoretical framework.

III. Rand's Commitment to Scientific Cooperation

Rand has an enviable history and reputation for cooperation with the scientific community. We see our interests being centered primarily in two areas, the government and academe. Our corporate mission is to bridge the gap between these two institutions. Our major motivation for joining the AIM community is to magnify our current and future impact in the areas of AI, cognitive psychology, and health services. We emphasize quality research and effective, scholarly communication. As prominent members of the ARPANET community, we have developed a number of systems that are widely employed (such as editors, production system packages, graphics, and message systems). Two of our staff members recently sponsored an AI workshop on Pattern-Directed Inference Systems, edited the proceedings published in SIGART Newsletter, and edited a forthcoming book published by Academic Press. The exportation and importation of software is a regular part of our business, and we look to joining the AIM community as an opportunity to expand our horizons and influence.

IV. Other Research in Information Processing Systems and Health

Briefly, Rand has recently completed development of the CLINFO system for clinical research and has now transferred it to NIH. NIH is currently soliciting proposals for a private enterprise to supply and support 30 mini-computer based installations of CLINFO. In addition to our Health Insurance Study, a number of other investigations are underway in health services, including the economics of health, the legal implications of computer systems, privacy of medical records, and biosystems modeling. Other current AI and cognitive sciences projects cover exemplary programming (creating heuristic agents by examples of their behavior), partial matching as a basis for concept discovery and rule induction, reasoning about maps and spatial information, and a hierarchical model of planning (based on the Hearsay-II model) that attempts to incorporate multiple, diverse knowledge sources cooperating to formulate a plan by integrating concepts from the highest levels (goals, policies) to the lowest levels (procedures, operations, and hypothetical assertions).

The Appendix provides a representative set of recent publications of our staff members and active consultants, some of whom are listed below:

Bob Anderson -- intelligent personal computers, graphics
Dave Drew -- rule-based systems for governmental administration
Bill Faught -- modeling human emotions and behavior, language
Stock Gaines -- secure operating systems
Gabe Groner -- computer aids for clinical research
Barbara Hayes-Roth -- knowledge organization, learning and planning
Rick Hayes-Roth -- knowledge acquisition, knowledge system design
Dave Holzman -- information systems for management
Stan Rosenschein -- AI, efficiency, and programming languages
Perry Thorndyke -- knowledge representation, comprehension
Don Waterman -- heuristic modeling, machine learning
Al Williams -- Program Director for Health Sciences

A partial list of consultants:

Ed Feigenbaum -- artificial intelligence
Doug Lenat -- heuristic reasoning, learning and discovery
Raj Reddy -- artificial intelligence
Bill Schwartz -- artificial intelligence in medicine

APPENDIX -- SELECTED REFERENCES

[R. H. Anderson76a]
Anderson, R. H., and Gillogly, J. J. The Rand intelligent terminal agent (RITA) as a network access aid. [AFIPS Proceedings], [45], 1976, 501-509.

[R. H. Anderson77b]
Anderson, R. H., Gallegos, M., Gillogly, J. J., Greenberg, R. B., and Villanueva, R. RITA Reference Manual, R-1808-ARPA, The Rand Corporation, Santa Monica, 1977.

[R. H. Anderson77a]
Anderson, R. H. The use of production systems in RITA to construct personal computer "agents." In [Proceedings of the Workshop on Pattern-Directed Inference Systems], SIGART Newsletter No. 63, 1977, 23-28.

[Faught75]
Faught, W. S. Affect as motivation for cognitive and conative processes. [Proceedings of the Fourth International Joint Conference on Artificial Intelligence], Tbilisi, USSR, 1975, 893-899.

[Faught77]
Faught, W. S., Colby, K. M., and Parkison, R. C. Inferences, affects, and intentions in a model of paranoia. [Cognitive Psychology], [9], 1977, 153-187.

[B. Hayes-Roth77a]
Hayes-Roth, B. The evolution of cognitive structures and processes. [Psychological Review], [84], 1977, 260-278.

[B. Hayes-Roth75]
Hayes-Roth, B. and Hayes-Roth, F. Plasticity in memorial networks. [Journal of Verbal Learning and Verbal Behavior], [14], 1975, 506-522.

[B. Hayes-Roth77b]
Hayes-Roth, B. and Hayes-Roth, F. Concept learning and the recognition and classification of exemplars. [Journal of Verbal Learning and Verbal Behavior], [16], 1977, 321-338.

[B. Hayes-Roth77c]
Hayes-Roth, B. and Hayes-Roth, F. The prominence of lexical information in memory representations of meaning. [Journal of Verbal Learning and Verbal Behavior], [16], 1977, 119-136.

[F. Hayes-Roth76a]
Hayes-Roth, F. Patterns of induction and associated knowledge acquisition algorithms. In [Pattern recognition and artificial intelligence], Chen, C.H. (ed.). Academic Press, New York, 1976.

[F. Hayes-Roth76c]
Hayes-Roth, F. Uniform representations of structured patterns and an algorithm for the induction of contingency-response rules. [Information and Control], [33], 1976, 87-116.

[F. Hayes-Roth77a]
Hayes-Roth, F. Learning by Example. In [Cognitive psychology and instruction], Glaser, R., Lesgold, A. & Fokkema, S. (eds.). Plenum, New York, 1977 (in press).

[F. Hayes-Roth77b]
Hayes-Roth, F. The role of partial and best matches in knowledge systems. In [Pattern-directed inference systems], Waterman, D. A. and Hayes-Roth, F. (eds.). Academic Press, New York, in press.

[F. Hayes-Roth77e]
Hayes-Roth F. and Lesser, V. Focus of attention in the Hearsay II speech understanding system. [Proceedings of the Fifth International Joint Conference on Artificial Intelligence], Cambridge, 1977, 27-35.

[F. Hayes-Roth77f]
Hayes-Roth, F. and McDermott, J. Knowledge acquisition from structural descriptions. [Proceedings of the Fifth International Joint Conference on Artificial Intelligence], Cambridge, 1977, 356-362.

[Lenat77d]
Lenat, D. and Harris, G. Designing a rule system that searches for scientific discoveries. In [Pattern-directed inference systems], Waterman, D. A. and Hayes-Roth, F. (eds.). Academic Press, New York, in press.

[McDonald77]
McDonald, D. and Hayes-Roth, F. Inferential searches of knowledge networks as an approach to extensible language understanding systems. In [Pattern-Directed Inference Systems], Waterman, D.A. and Hayes-Roth, F. (eds.). Academic Press, New York, in press.

[Mitchell76]
Mitchell, B. M., and Schwartz, W. B. The financing of national health insurance. Rand R-1711-HEW. 1976.

[Newell57]
Newell, A., Shaw, J., and Simon, H. Empirical explorations of the logic theory machine: A case study in heuristics. Report P-951, The Rand Corporation, Santa Monica, 1957.

[Newhouse74]
Newhouse, J. P. The Health Insurance Study--a summary. Rand R-965-1-OEO. 1974.

[Parkison77]

Parkison, R. C., Colby, K. M., and Faught, W. S. Conversational language comprehension using integrated pattern-matching and parsing. [Artificial Intelligence], (in press).

[Palley74]
Palley, N. A., and Groner, G. F. A survey of clinical investigators and their information processing activities. Rand R-1539-NIH. 1974.

[Palley76]
Palley, N. A., Groner, G. F., Sibley, W. L., and Hopwood, M. D. CLINFO user's guide: release one. Rand R-1543-1-NIH. 1976.

[Rosenschein75a]
Rosenschein, S. J. Structuring a pattern space, with applications to lexical information and event interpretation. Ph.D. dissertation, University of Pennsylvania, Philadelphia, 1975.

[Rosenschein75b]
Rosenschein, S. J. How does a system know when to stop inferencing. [American Journal of Computational Linguistics], Microfiche 36, 1975.

[Rosenschein77a]
Rosenschein, S. J. The production system: Architecture and abstraction. In [Pattern-directed inference systems], Waterman, D. A. and Hayes-Roth, F. (eds.). Academic Press, New York, in press.

[Thorndyke77a]
Thorndyke, P. W. Cognitive structures in comprehension and memory of narrative discourse. [Cognitive Psychology], [9], 1977, 77-110.

[Thorndyke76]
Thorndyke, P. The role of inferences in discourse comprehension. [Journal of Verbal Learning and Verbal Behavior], [15], 1976, 437-446.

[Thorndyke77b]
Thorndyke, P. Knowledge transfer in learning from texts. In [Cognitive psychology and instruction], Glaser, R., Lesgold, A., Fokkema, S. (eds.). Plenum Publishers, Amsterdam, 1977.

[Waterman70]
Waterman, D. A. Generalization learning techniques for automating the learning of heuristics. [Artificial Intelligence], [1], 1970, 121-170.

[Waterman75]
Waterman, D. A. Adaptive production systems. [Proceedings of the Fourth International Joint Conference on Artificial Intelligence], Tbilisi, USSR, 1975, 296-303.

[Waterman77a]
Waterman, D. A. A rule-based approach to knowledge acquisition for man-machine interface programs. P-5895, Rand Corporation, Santa Monica, 1977.

[Waterman77b]
Waterman, D. A. Rule-directed interactive transaction agents: An approach to knowledge acquisition. R-2171-ARPA, Rand Corporation, Santa Monica, 1977.

[Waterman77c]
Waterman, D. A. and Jenkins, B. Heuristic modeling using rule-based computer systems. P-5811, Rand Corporation, Santa Monica, 1977.