

believe the resulting "pure" models would be too complex to be useful. In combining the two approaches, concepts best modeled in a physiological paradigm can be expressed within the mathematical model while concepts best modeled symbolically can be represented within the temporal network. We are continuing to explore this architecture in the VentPlan system designed to assist with data interpretation in the Intensive Care Unit application area.

C.2.2 Documentation

In 1986, we videotaped a lecture and demonstration of the ONCOCIN and OPAL systems at the XEROX Palo Alto Research Center. This videotape is available for loan from our offices. Our previous videotapes have been shown at scientific meetings and have been distributed to many researchers in other countries. The publications described below further document our recent work on ONCOCIN. We have sent copies of our system to the University of Pittsburgh, and have distributed it to the National Library of Medicine. We have developed a user manual, description of sample interaction, reference card, and graphical flowchart to help with training in the use of ONCOCIN.

D. Publications Since January, 1988

- 1) Musen, M.A. Generation of Model-Based Knowledge-Acquisition Tools for Clinical-Trial Advice Systems. KSL-88-06. Doctoral dissertation, Medical Information Sciences Program, Medical Computer Science Group, Stanford University, January 1988 Updated version appears as Automated Generation of Model-Based Knowledge-Acquisition Tools. London: Pitman, 1989.
- 2) Wulfman, C.E., Isaacs, E.A., Webber, B.L., and Fagan, L.M. Integration discontinuity: Interfacing users and systems. Memo KSL-88-12, February 1988. Proceedings of Architectures for Intelligent Interfaces: Elements and Prototypes, pp. 57-68, Monterey CA, March 29-April 1, 1988.
- 3) Musen, M.A. Conceptual models of interactive knowledge-acquisition tools. Report KSL-88-16, March 1988. Proceedings of European Knowledge Acquisition Workshop (EKAW'88), Bonn, FRG, June 1988. Gesellschaft für Mathematik und Datenverarbeitung (GMD) Technical Report #143. Knowledge Acquisition, in press.
- 4) Musen, M.A. Generation of knowledge-acquisition tools from clinical-trial models. Report KSL-88-26, March 1988. Lecture Notes in Medical Informatics (P.L. Reichertz and D.A.B. Lindberg, eds.), Vol. 35, pp. 630-635. Proceedings (Medical Informatics Europe '88), Oslo, Norway: Springer-Verlag, 17-20 August 1988.

- 5) Musen, M.A. An editor for the conceptual models of interactive knowledge-acquisition tools. Report KSL-88-44, June 1988. Proceedings of the Third Workshop on Knowledge Acquisition for Knowledge-Based Systems, Banff, Alberta, Canada, November 1988. Also to appear in the International Journal of Man-Machine Studies, in press. A shorter version of this paper was published in the Knowledge Acquisition Special Issue of SIGART Newsletter, 108:45-55, April 1989.
- 6) Musen, M.A. and van der Lei, J. Of brittleness and bottlenecks: Challenges in the creation of pattern-recognition and expert-system models. Report KSL-88-59, August 1988. Pattern Recognition and Artificial Intelligence: Towards an Integration of Techniques (E.S. Gelsema and L.N. Kanal, eds.), pp. 335-352. New York: Elsevier, 1988.
- 7) Tu, S.W., Kahn, M.G., Musen, M.A., Ferguson, J.C., Shortliffe, E.H. and Fagan, L.M. Episodic monitoring of time-oriented data for heuristic skeletal-plan refinement. Report KSL-87-70, August 1988. (submitted to CACM).
- 8) Langlotz, C.P. and Shortliffe, E.H. An analysis of categorical and quantitative methods for planning under uncertainty. Report KSL-88-63, September 1988. Proceedings of the Twelfth Annual Symposium on Computer Applications in Medical Care, pp. 114-118, Washington, D.C., November 1988.
- 9) Combs, D. The OPAL knowledge-acquisition tool: Functional specification. Technical documentation, internal working memo.
- 10) Combs, D. ODIE: A system for design and management of form-based interfaces. Technical documentation, internal working memo.
- 11) Shwe, M., Tu, S.W. and Fagan, L.M. Validating the knowledge base of a therapy-planning system. Report KSL-88-70. October 1988. Methods of Information in Medicine, 28(1):36-50, 1989. Musen, M.A. Generation of visual languages for development of knowledge-based systems. Report KSL-88-73, December 1988. To appear as a chapter in Visual Languages, Volume II (R.R. Korfhage, E. Jungert, and T. Ichikawa, eds.), New York: Plenum, 1989.
- 12) Shortliffe, E.H. Testing reality: The introduction of decision-support technologies for physicians. Report KSL-88-82, December 1988. Published as an editorial in Methods of Information in Medicine, 28(1):1-5, 1989.
- 13) Beinlich, I.A., Suermondt, H.J., Chavez, R.M., and Cooper, G.F. The ALARM monitoring system: A case study with two probabilistic inference techniques for belief networks. Report KSL-88-84, January 1989. To appear in AI in Medicine, London 1989.
- 14) Rutledge, G., Thomsen, G., Beinlich, I., Farr, B., Kahn, M., Sheiner, L., and Fagan, L. VentPlan: An architecture for combining qualitative and quantitative computation. Report KSL-89-04, January 1989.

- 15) Musen, M.A. Languages for knowledge acquisition: Building and extending models. Report KSL-89-07, January 1989. Proceedings of the AAAI Spring Symposium on Knowledge System Development Tools and Languages. American Association for Artificial Intelligence, Stanford CA, March 1989.
- 16) Musen, M.A. Widening the knowledge-acquisition bottleneck: Automated tools for building and extending clinical models. Report KSL-89-09, February 1989. Proceedings of the AAMSI Congress '89, 1989, pp 2-7.
- 17) Farr, B.R. Decision-theoretic evaluation of therapy plans. Report KSL-89-20. Submitted to the Proceedings of the Thirteenth Annual Symposium on Computer Applications in Medical Care, March 1989.
- 18) Reed, J. Building decision models that modify decision systems. Report KSL-89-21. Submitted to the Proceedings of the Thirteenth Annual Symposium on Computer Applications in Medical Care, March 1989.
- 19) Musen, M.A. Automated support for building and extending expert models. Report KSL-89-26, March 1989. To appear in a special issue of Machine Learning.
- 20) Musen, M.S. and van der Lei, J. Knowledge engineering for clinical consultation programs: Modeling the application area. Report KSL-89-28, March 1989. Methods of Information in Medicine, 28(1):28-35, 1989.
- 21) Kahn, M.G., Fagan, L.M., and Sheiner, L.B. Model-based interpretation of time-varying medical data. Report KSL-89-34, April 1989.
- 22) Langlotz, C.P. and Shortliffe, E.H. The critiquing approach to automated advice and explanation: Rationale and examples. KSL-89-46, May 1989. To appear in Expert Knowledge and Explanation: The Knowledge-Language Interface (C. Ellis, ed.), London: Ellis Horwood, 1989.

E. Funding Support

Grant Title: "Therapy-planning strategies for consultation by computer"

Principal Investigator: Edward H. Shortliffe

Project Management: Lawrence M. Fagan

Agency: National Library of Medicine

ID Number: LM-04136

Term: April 1987 to March 1990

Total award: \$380,123

Grant Title: Postdoctoral Training in Medical Information Science

Principal Investigator: Edward H. Shortliffe

Project Management: Edward H. Shortliffe

Agency: National Library of Medicine

ID Number: 1 T32 LM07033

Term: July 1, 1984 - June 30, 1989

Total award: \$903,718

II. INTERACTIONS WITH THE SUMEX-AIM RESOURCE

A. Medical Collaborations and Program Dissemination via SUMEX

A great deal of interest in ONCOCIN has been shown by the medical, computer science, and lay communities. We are frequently asked to demonstrate the program to Stanford visitors. We also demonstrated our developing workstation code in the Xerox exhibit in the trade show associated with AAAI-84 in Austin, Texas, IJCAI-85 in Los Angeles, AAAI-86 in Philadelphia, and Medinfo 86. Physicians have generally been enthusiastic about ONCOCIN's potential. The interest of the lay community is reflected in the frequent requests for magazine interviews and television coverage of the work. Articles about MYCIN and ONCOCIN have appeared in such diverse publications as *Time* and *Fortune*, and ONCOCIN has been featured on the "NBC Nightly News," the PBS "Health Notes" series, and "The MacNeil-Lehrer Report." Most recently it appeared in a special on Artificial Intelligence for TV Ontario (Canadian PBS station) and "Physician's Journal Update" on the Lifetime Cable Network. Due to the frequent requests for ONCOCIN demonstrations, we have produced a videotape about the ONCOCIN research which includes demonstrations ONCOCIN and OPAL. The tape has also been shown to both national and international researchers in biomedical computing. We are producing a tape that describes our speech research.

Our group also continues to oversee the MYCIN program (not an active research project since 1978) and the EMYCIN program. Both systems continue to be in demand as demonstrations of expert systems technology. MYCIN has been demonstrated via networks at both national and international meetings in the past, and several medical school and computer science teachers continue to use the program in their computer science or medical computing courses. Researchers who visit our laboratory often begin their introduction by experimenting with the MYCIN/EMYCIN systems. We also have made the MYCIN program available to researchers around the world who access SUMEX using the GUEST account. EMYCIN has been made available to interested researchers developing expert systems who access SUMEX via the CONSULT account. One such consultation system for psychopharmacological treatment of depression, called Blue-Box (developed by two French medical students, Benoit Mulsant and David Servan-Schreiber), was reported in July of 1983 in *Computers and Biomedical Research*. The EMYCIN experience is now well disseminated via commercial products.

B. Sharing and Interaction with Other SUMEX-AIM Projects

The community created on the SUMEX resource has other benefits which go beyond actual shared computing. Because we are able to experiment with other developing systems, such as INTERNIST/CADUCEUS, and because we frequently interact with other workers (at AIM Workshops or at other

meetings), many of us have found the scientific exchange and stimulation to be heightened. Several of us have visited workers at other sites, sometimes for extended periods, in order to pursue further issues which have arisen through SUMEX- or workshop-based interactions. In this regard, the ability to exchange messages with other workers, both on SUMEX and at other sites, has been crucial to rapid and efficient dissemination of ideas. Certainly it is unusual for a small community of researchers with similar scholarly interests to have at their disposal such powerful and efficient communication mechanisms, even among those researchers on opposite coasts of the country.

During this past three years, we have had extensive interactions with Randy Miller at Pittsburgh. Via floppy disks and SUMEX, we have experimented with several versions of the QMR program. The interaction was very much facilitated by the availability of SUMEX for communication and data transmission. Several recent papers have been written to describe collaborations between students in our training program and the group at the University of Pittsburgh and Carnegie-Mellon University.

C. Critique of Resource Management

Our community of researchers has been extremely fortunate to work on a facility that has continued to maintain the high standards that we have praised in the past. The staff members are always helpful and friendly, and work as diligently to please the SUMEX community as to please themselves. As a result, the computer is as accessible and easy-to-use as they can make it. More importantly, it is a reliable and convenient research tool. We extend special thanks to Tom Rindfleisch for maintaining such high professional standards. As our computing needs grow, we have increased our dependence on special SUMEX skills such as networking and communication protocols. As described above, we are moving our software development to a combination of Lisp processor boards and Mac II's, and will need increased networking capabilities in this environment.

III. RESEARCH PLANS

A. Project Goals and Plans

In the coming year, there are several areas in which we expect to expend our efforts on the ONCOCIN System:

- 1) To generalize the reasoning and interaction components of the ONCOCIN system for other applications.
- 2) Extensions and generalizations of the OPAL and PROTÉGÉ knowledge acquisition experiments
- 3) Extensions to the strategic planning framework including research on temporal representations, mathematical modeling and summarization
- 4) To determine and begin implementation of ONCOCIN on more accessible hardware platforms..

B. Justification and Requirements for Continued SUMEX Use

All of our research takes place in the context of the SUMEX resource. The development of our project will continue to take place on LISP machines which we have purchased or which have been donated by the XEROX Corporation, with a gradual transition to Mac II's. The word processing and communication requirements will be met by the Mac II's plus the excellent network services provided by SUMEX.

C. Requirements for Additional Computing Resources

Most of our CPU needs are met with our current equipment. However, our requirements for high bandwidth communication facilities are increased as we have an increasingly distributed environment. We would appreciate support that provides equivalent resources to what we are used to on the DEC-20 including mail support, and file servers. We will continue to need access to the international networks that we now use for much of our communication with colleagues. For example, Mark Musen (who recently finished his Ph.D., and is now an Assistant Professor in our department) was in Holland for 9 months. During this period, we were in constant electronic communication with him over a set of networks. This has been a significant asset to our research project.

D. Recommendations for Future Community and Resource Development

The continuation of network support and file service are our major needs. Help with software selection and implementation for our Mac II's is also an important requirement to maintain research continuity in our heterogeneous environment. Maintaining the excellent mail service and access to the international networks is also essential for our research.

IV.A.4. PENGUIN Project

The PENGUIN Project: Toward Expert Database Systems in Biomedicine

Gio Wiederhold, Ph.D.
Departments of Computer Science and Medicine
Stanford University

Thierry Barsalou, M.D.
Medical Computer Science
Stanford University

I. SUMMARY OF RESEARCH PROGRAM

Crucial problems encountered in the management and structuring of large quantities of knowledge are being addressed by the KBMS/KSYS project, primarily under DARPA sponsorship. The RADIX project (phased out last year) on automated knowledge acquisition from large medical databases has provided valuable experience for the KBMS/KSYS group. The KBMS/KSYS work, in turn, is leading to another important medical application: PENGUIN, which investigates combining knowledge-base and database technology into so-called expert database systems. This project was initially supported by KBMS, but is now fully funded by a grant from NLM.

A. PENGUIN Project — Rationale

Databases and expert systems share a common goal — generating useful information for action — but accomplish their tasks separately, using different principles. It is clear, however, that future information systems will require both the problem-solving capabilities of expert systems (ES's) and the data-handling capabilities of database management systems (DBMS's). Indeed, combining database and expert system technologies into expert database systems (EDS's) is an emerging research area. One can define an EDS as "a system for developing applications requiring knowledge-directed processing of shared information". From a perspective of developing advanced biomedical information systems, this definition conveys two precise scenarios: (1) enhancing DBMS's with structuring and manipulation tools that take more semantics into account; (2) allowing ES's to access and to handle efficiently information stored in database(s).

The object-oriented paradigm has gained much attention in recent years. In the database field, object-oriented DBMS's have emerged. The concept of an entity is also widely used by database design tools. In the field of artificial intelligence, frames are a well-known knowledge representation scheme. Although frames were conceived separately from the object paradigm, the two are in fact consistent with each other. In this research project, called PENGUIN, we investigate the hypothesis that the object-oriented approach can also serve as a unifying scheme for developing EDS's.

We have the opportunity to explore this hypothesis in a practical biomedical environment. Fluorescence Activated Cell Sorting (FACS) is emerging as a major source of information for biomedical research and clinical practice. Currently, achieving good FACS performance requires analyzing and integrating various complex data and knowledge sources. PENGUIN is thus aimed at developing methods for bringing together expert system and database technologies in an integrated advice system that could fulfill the information needs of FACS investigators. Central to this work is a very close collaboration between researchers in the Medical Information Sciences Program and the Departments of Computer Science and Genetics.

B. PENGUIN — Medical Relevance and Collaboration

FACS has already demonstrated significant promise in clinical as well as research-oriented areas. As indicated earlier, basic studies identifying lymphocyte subpopulations have now been translated into clinical applications including monitoring of AIDS patients. Similar FACS applications can now be found in such diverse specialties as Hematology, Oncology, Infectious Diseases and even Gynecology-Obstetrics where current studies are evaluating FACS analysis of maternal blood sample as a potential non-invasive method for prenatal diagnoses.

The major block to proliferation of these kinds of valuable studies involves the requirement for greater skills in reagent selection and FACS machine operation than are typically available in basic and clinical research settings.

Thus developing sophisticated computer tools that provide a new level of automatic analysis and control for FACS and greatly facilitates FACS use by reducing the need for on-site human expertise should significantly improve the potential for using this versatile methodology in biomedical research and clinical practice.

C. PENGUIN — Highlights of Research Progress

Our current effort focuses on developing a computer-based advisory system to assist the FACS investigator in designing experiment protocols. As mentioned above, this system will combine database and artificial intelligence methodologies in an integrated framework, as to provide 1) several levels of abstraction for information management and retrieval from a relational database system, 2) access to and integration of the results of past similar experiments as additional units of information through an interface with existing databases of FACS data and 3) inference capabilities coupled to the database for high-level interactions with scientists during the design process. Although this work is motivated by a specific application, the formal design of the system will be domain-independent; it is then our belief that ideas, principles and programs developed in this process will be applicable to other medical and non-medical areas.

In the past two years, we have developed the concept of an object-based interface on top of a relational database system and implemented an initial

prototype of the interface. We have drawn an analogy between the notions of object and database view. Using this analogy, we have defined three components in the object interface:

- 1) The object *generator* maps relations into object templates where each template can be a complex combination of join (combining two relations through shared attributes) and projection (restricting the set of attributes of a relation) operations on the base relations. In addition, an object network groups together related templates, thereby identifying different object views of the same database. The whole process is knowledge-driven, using the semantics of the database structure. We define the object schema as the set of object networks constructed over a given database. Like the data schema for a relational database, the object schema represents the domain-specific information needed to gain access to PENGUIN's objects; this information enables us to combine well-organized, regular tabular structures — the relations — into complex, heterogeneous entities — the objects.
- 2) The object *instantiator* provides nonprocedural access to the actual object instances. First, a declarative query (e.g., Select instances of template x where attribute y < 0.5) specifies the template of interest. Combining the database-access function (stored in the template), and the specific selection criteria, PENGUIN automatically generates the relational query and transmits it to the DBMS, which in turn transmits back the set of matching relational tuples. In addition to performing the database-access function, the object template specifies the structure and linkage of the data elements within the object. This information is necessary for the tuples to be correctly assembled into the desired instances. Those instances are then made available to the expert system directly, or to the user through a graphic interface.
- 3) The object *decomposer* implements the inverse function; that is, it maps the object instances back to the base relations. This component is invoked when changes to some object instances (e.g., deletion of an instance, update of some attributes) need to be made persistent at the database level. An object instance is generated by collapsing (potentially) many tuples from several relations. By the same token, one update operation on an object may result in a number of update operations that need to be performed on the base relations.

In addition to augmenting the relational model with an object layer, we have devoted a significant portion of our effort to user-interface issues, as one of the main challenges faced by designers of information systems is the development of high-grade interfaces. Indeed, wide acceptance and use of any computer program presuppose a good interface. Our work follows from two basic principles. First, domain-specific interfaces should be separated from all other aspects of the system, which could then offer a number of interfaces tailored to the problem at hand. Second, some intelligence should be built into the interfaces to make them truly cooperative, active agents

instead of simply passive media. Our research hypothesis is that a hypertext approach can facilitate the design of such intuitive user interfaces for decision-support systems. We are exploiting the HyperCard authoring tool for that purpose.

Using HyperCard, one can design, prototype, and debug a tailored interface in a single session—interface that takes full advantage of a visual representation of the underlying database (through, for instance, icons, buttons, and graphics). Within the taxonomy of direct-manipulation interfaces, PENGUIN employs visual correlation to specify objects and actions. Inasmuch as visual-correlation methods require the user to recognize rather than to recall the appropriate command name, they typically place a lesser burden on human memory than command languages do. Most important, perhaps, refinement of the interface can occur without writing a line of code through the large palette of tools that we have assembled.

Preliminary results of PENGUIN's implementation indicate that an object-based architecture provides (1) efficient access to and manipulation of complex units of biomedical information while preserving the advantages associated with persistent storage of data in relational format, and (2) a domain-independent, bidirectional communication channel between relational database systems and expert systems.

The latter result is based on our ongoing work toward integrating PENGUIN and KNET. Developed in our laboratory by R. Martin Chavez, KNET is a general-purpose environment for constructing probabilistic, knowledge-intensive systems based on belief networks and decision networks. Integration of PENGUIN and KNET requires translating KNET knowledge bases into relational format. For that purpose, we have defined a general relational schema that can accommodate any domain-specific KNET knowledge base. Thanks to that general schema, the translation process is for the most part automated. We now need to establish communication between the two systems. Because KNET is object-oriented, PENGUIN's object layer represents a natural communication channel. We therefore envision a tripartite architecture with a database layer and a knowledge-based layer connected by the object layer.

Finally and in parallel to system development, we are currently engaged in a knowledge acquisition process to define the structure of the database and of the knowledge base. Through interviews with our experts in the Genetics department, we are eliciting the structure of the various components (genetic, histological, and serological information) of the FACS reagents database.

D. Publications of the PENGUIN project

- 1) Barsalou, T. and Wiederhold G.: "A cooperative hypertext interface to relational databases"; submitted to the Thirteenth Annual Symposium on Computer Application in Medical Care, March 1989.

- 2) Barsalou, T., Chavez, R.M. and Wiederhold G.: "Hypertext interfaces for decision-support systems: A case study"; To appear in the Proceedings of MEDINFO '89, IFIP, Beijing, China, October 1989.
- 3) Barsalou, T. and Wiederhold G.: "Knowledge-based mapping of relations into objects"; To appear in Computer Aided Design, 1989.
- 4) Barsalou, T.: "An object-based architecture for biomedical expert database systems"; in R.A. Greenes (editor), Proceedings of the Twelfth Symposium on Computer Applications in Medical Care, pages 572—578, IEEE Computer Society Press, Washington, D.C., November 1988.
- 5) Barsalou, T. and Wiederhold G.: "Applying a Semantic Model to an Immunology Database"; in W.W. Stead (editor), Proceedings of the Eleventh Symposium on Computer Applications in Medical Care, pages 871-877, IEEE Computer Society Press, Washington, D.C., November 1987.
- 6) Barsalou, T., Moore, W.A., Herzenberg, L.A. and Wiederhold G.: "A Database System to Facilitate the Design of FACS Experiment Protocols" (abstract); Cytometry, Vol.97, August 1987.
- 7) Law, K.H. and Barsalou T.: "Applying a semantic structural model for engineering design"; To appear in the Proceedings of the Computers in Engineering Conference, American Society of Mechanical Engineers, Anaheim, CA, August 1989.
- 8) Wiederhold, G., Rathmann, P.K., Barsalou, T., Lee, B.S. and Quass D.: "Partitioning and composing knowledge"; Submitted to Information Systems, 1989.
- 9) Wiederhold, G., Barsalou, T. and Chaudhuri S.: "Managing objects in a relational framework"; Technical report No. STAN-CS-89-1245, Computer Science Department, Stanford University, January 1989.
- 10) Wiederhold, G.: "Hospital Information Systems"; in Encyclopedia of Medical Devices and Instrumentation, vol.3, Webster, John G.(ed), Wiley 1988, pp.1517—1542.
- 11) Wiederhold, G.: "Sharing and partitioning large knowledge bases"; in Appelrath, Cremers, and Schiltknecht (eds), 'PROTOS, Prolog tools for building expert systems', EUREKA report EU56, Basel, Switzerland, pp. 77—83, December 1988.
- 12) Wiederhold, G.: *File Organization for Database Design* McGraw-Hill Book Company, 1988.

- 13) Wiederhold, G., Michael G. Walker, Waqar Hasan, Surajit Chaudhuri, Arun Swami, Sang K. Cha, XiaoLei Qian, Marianne Winslett, Linda DeMichiel, and Peter K. Rathmann: "KSYS: An Architecture for Integrating Databases and Knowledge Bases"; Computer Science Department, Stanford University, May 1987; in Amar Gupta and Stuart Madnick (editors) 'Technical Opinions Regarding Knowledge- Based Integrated Information Systems Engineering', MIT, 1987,
- 14) Wiederhold, G.: "Knowledge versus Data"; Chapter 8 of 'On Knowledge Base Management Systems: Integrating Artificial Intelligence and Database Technologies' (Brodie, Mylopoulos, and Schmidt, eds.), Springer Verlag, June 1986, pages 77 to 82.
- 15) Wiederhold, G., Blum, R.L., and Walker M.: "An Integration of Knowledge and Data Representation"; Proc. of Islamorada Workshop, Feb.1985, Computer Corporation of America, Cambridge MA; Chapter 29 of 'On Knowledge Base Management Systems: Integrating Artificial Intelligence and Database Technologies' (Brodie, Mylopoulos, and Schmidt, eds.), Springer Verlag, June 1986, pages 431 to 444.
- 16) Wiederhold, G.: *Views, Objects, and Databases*. IEEE Computer 19(12):37-44, December, 1986.
- 17) Missikoff, M., and Wiederhold G.: *Towards a Unified Approach for Expert and Database Systems*. In 'Expert Database Systems', Larry Kerschberg (editor), Benjamin/Cummings, 1986, pages 383-399; also in Proceedings of First Workshop on Expert Database Systems, Kiawah Island, South Carolina, Oct. 1984, vol.1, pp.186-206.
- 18) Wiederhold, G.: *Knowledge Bases*; Future Generations Computer Systems, North-Holland, vol.1 no.4; April 1985, pp.223—235.
- 19) Wiederhold, G.: *Database Design* (in the Computer Science Series); McGraw-Hill Book Company, New York, NY, May 1977, 678 pp. Second edition, Jan. 1983, 768 pp.
- 20) Wiederhold, G.: In D.A.B. Lindberg and P.L. Reichertz (Eds.), *Databases for Health Care*, Lecture Notes in Medical Informatics, Springer-Verlag, 1981.
- 21) Wiederhold, G.: "Database technology in health care"; J. Medical Systems 5(3):175-196, 1981.

E. Funding Support Status

1987-1990 Principal Investigator: Gio Wiederhold 25% effort
 FACS-Penguin: An Expert Workstation for Flow Cytometry
 (NLM/NIH : \$856,449)

1987-1991 Principal Investigator: Gio Wiederhold 10%
 Support for Parallel Design in an Engineering Information System
 (NSF DTMP 8619595 \$153,766, \$146,360 prop.,\$156,044 prop.)

1987-1988 Principal Investigator: Gio Wiederhold 5%
 Validation of Knowledge from a Database
 (IBM KBS Menlo Park, \$50,051)

1986-1990 Principal Investigator: Gio Wiederhold 50% effort
 Knowledge Base Management Structures
 (ONR/SPAWAR/ARPA,N39-84-C-211, task 7: \$1,756,410)

1987-1988 Principal Investigator: Gio Wiederhold 5%
 Reasoning about R1ME
 (Digital Equipment Corp: \$131,337)

1987-1990 Associate Investigator: Gio Wiederhold 10% effort
 Integrating Knowledge and DBMS
 (SRI AI/ONR/SPAWAR/ARPA: \$35,003 Stanford Subcontract)

1988-1990 Principal Investigator: Gio Wiederhold 10% effort
 PARADATA: Databases on Parallel Computers
 (NSF/IRI/K&DSP, year 1: \$148,116, expected year 2 \$159,689)

1988 Principal Investigator 10% effort
 DADAISM: DBMS in and for ADA-supported Information System
 Management
 (SRI/NRL/STARS/DoD: supplement \$30,000)

1988 Principal Investigator 20% summer, 10% academic year
 (6 months) DADAISM: DBMS in and for ADA-supported Information
 System Management
 (SRI/NRL/STARS/DoD: \$190,861)

1988-1991 Principal Investigator: 5% effort
 Computer-Assisted Analysis of Auroral Images Obtained from
 High Altitude Polar Satellites
 (CESDIS/USRA/NASA, approved,\$577,254 requested, being funded at
 approx. \$500,000, with supplementary support for Dr. Robert Clauer)

1988 Co-investigator with Kincho Law 5% effort
 Object Databases for CAD/CAE data
 (Stanford/CIFE, \$34,500)

1990-1994 Associate Investigator, Paul Losleben (PI) 10% + 10% effort
 Micro Factory Project, Database for Manufacturing Expert,
 Systems Process Design Database,
 (Texas Instruments/NSF, \$287,442 + \$288,877 out of \$4,540,547)

Requests in process

1989-1994 Consultant Investigator, James Fries (PI) 5% effort
 ATHOS: AIDS- Time-oriented Health Outcome Study;
 (NCHSR/PHS/HEW, \$7,562,601 requested)

II. INTERACTIONS WITH THE SUMEX-AIM RESOURCE

A. Collaborations

SUMEX AIM provides the central communication node for our research. While individual experiments tend to take place on workstations, the availability of SUMEX to load and edit files, communicate with colleagues and peers is absolutely critical.

B. Interactions with Other SUMEX-AIM Projects

During the current reporting year we have had frequent interaction with members of other SUMEX projects; for example, presentation of research results at Stanford Medical Information Science Colloquia, discussions of automated discovery and automated summarization, practical programming issues, and training of Medical Computer Science Students in the use of KEE, Lisp workstations, and so on. The SUMEX community is an invaluable resource for providing such interaction.

C. Critique of Resource Management

The new Sun-4 system provides acceptable performance. The archiving facility (SAFE) seems to be very loaded, so that we are forced to make decisions on retention of past material more frequently than optimal. As users of large databases this is one of the resources upon we must draw frequently.

The SUMEX resource management continues to be accessible and quite helpful. New networking links, for instance to the Genetics Department equipment, make new collaborations viable.

III. RESEARCH PLANS

A. Project Goals and Plans

The long-range goal of the PENGUIN project is to integrate our experiment design advisory system with the set of programs developed in the Genetics department for FACS operation and control, thereby defining a comprehensive "FACS Workstation" that should considerably facilitate effective utilization of FACS technology. In the shorter term, we will:

- 1) complete the development and testing of the object interface. The interface should then be made a fully domain-independent module that could be applied in different contexts. One such context is civil engineering where the concepts developed in PENGUIN are being applied to engineering design problems.
- 2) integrate the PENGUIN and KNET systems, so that PENGUIN can act as an intelligent data server for KNET.

- 3) validate the structure of the database and make it available to the scientists in the Genetics department, for further refinement as well as for routine data storage and retrieval.
- 4) move fully toward a distributed environment with Macintoshes as local workstations connected to a central VAX computer, acting as a database server.
- 5) develop the expert system module which will be coupled to the database through the object interface. The expert system module is aimed at assisting people in an experimental design task. This can actually be done by taking two different approaches, one being the actual automatic design of the experiment (i.e., a planning task), the other being the critique of a protocol proposed by the user. Although we recognize both directions as important, we will focus on the first approach because knowing how to properly plan new experiments is a necessary step prior to any critiquing process. Although the problem solving strategies of our experts are not yet fully determined, we plan to investigate the general concept of hierarchical planning.

B. Justification and Requirements for Use of SUMEX

The PENGUIN project is dependent on new methods of communication among clinical and academic researchers.

C. Recommendations for Resource Development

The movement toward workstations is obviously the direction of the future. We hope that the communication capability, now supported via the central SUMEX facility will continue to be supported. We continue to make constant use of this communication with our colleagues within and outside of Stanford, preparation of large documents, file and database handling, and program demonstrations.

IV.A.5. PROTEAN Project

Oleg Jardetzky
Nuclear Magnetic Resonance Lab, School of Medicine
Stanford University

Bruce Buchanan, Ph.D.
Computer Science Department
Stanford University

I. SUMMARY OF RESEARCH PROGRAM

A. Project Rationale

The goals of this project have been related both to biochemistry and artificial intelligence: (a) use existing AI methods to aid in the determination of the 3-dimensional structure of proteins in solution, and (b) use protein structure determination as a test problem for experiments with large scale constraint satisfaction, an area of increasing interest to artificial intelligence. Empirical data from nuclear magnetic resonance (NMR) and other sources may provide enough constraints on structural descriptions to allow protein chemists to bypass the laborious methods of crystallizing a protein and using X-ray crystallography to determine its structure. This problem exhibits considerable computational complexity, but by formulating it as search problem, it should be possible to utilize many of the knowledge-based techniques developed in AI for dealing with large search spaces.

B. Medical Relevance

The molecular structure of proteins is essential for understanding many problems of medicine at the molecular level, such as the mechanisms of drug action. Using NMR data from proteins in solution will allow the study of proteins whose structure cannot be determined with other techniques, and will decrease the time needed for the determination.

C. Highlights of Progress

With the departure of Prof. Bruce Buchanan to the University of Pittsburgh, work on the PROTEAN project is nearing conclusion and has focussed on the completion of thesis work by Bruce Duncan and Russ Altman.

The core strategy of this project for assembling protein structures involves functionality at the atomic as well as the more abstract solid level, and demonstrates an application of our general approach, which we call *heuristic refinement*.

The heuristic refinement method falls within an *exclusion paradigm* for interpreting protein structure. In this paradigm all the atoms in the protein are initially assumed to be everywhere with respect to some arbitrary coordinate system. As constraints are introduced these initially infinite *accessible volumes* are gradually reduced until the remaining accessible

volumes are small enough so that a representative set of structures may be enumerated.

The exclusion paradigm has not been tried previously because it appears to be computationally intractable for any reasonably sized protein. For this reason most current techniques (non AI-based) use an *adjustment paradigm* to optimize a single structure towards a global minimum. These techniques all share difficulties common to optimization, namely lack of convergence for large numbers of variables or for starting structures far from the global minimum. In addition, for reasons of computation time, they are not able to provide a representative sample of all structures compatible with the constraints.

The exclusion paradigm, by systematically eliminating structures incompatible with the constraints, has the advantage that all, or at least a representative set, of structures compatible with the constraints are retained. In order to deal with the combinatorics the heuristic refinement method formulates the problem as one of search, which is the fundamental paradigm of artificial intelligence. In particular, the problem is formulated as a geometric constraint satisfaction problem (GCSP), which can be solved by backtrack search if the size of the atomic accessible volumes can be made small enough.

The complexity of solving a GCSP depends on the number of objects and the number of possible locations that each object can have in space. Therefore, the heuristic refinement method utilizes four techniques to reduce these two numbers to the point that backtrack search may be applied. The effect of each of these techniques is to utilize knowledge of protein structure and of techniques for solving GCSP's to prune the search tree:

- 1) **Problem decomposition.** The protein is broken into logical subparts such as secondary structures and sidechains. Each of these is treated as a separate GCSP, which is partially solved. Solutions to the subproblems either determine constraints at more abstract levels or are combined into larger subproblems.
- 2) **Problem abstraction.** Groups of locally highly constrained atoms are represented as single abstract "solid" level objects. Solutions to atomic level GCSP's are used to create abstract "solid" level constraints, resulting in a solid level GCSP with far fewer objects. Solutions at the solid level are then expanded to the atomic level, resulting in fewer possible locations for the atoms than would have been present if the initial solid level processing had not occurred. Almost all the work reported in previous reports was at this abstract solid level, but it is the current refinement to the atomic level that is of most interest to biochemists.
- 3) **Local satisfaction of constraints.** In each local GCSP filtering operations called network consistency algorithms are applied, allowing the accessible volumes to be reduced by looking at constraints pairwise (or to higher

order) rather than all at once. These operations, which were initially developed for computer vision, are becoming more popular in AI because of their utility in many forms of constraint satisfaction problems.

- 4) Heuristic control. At each point in the problem solving there are many possible constraint satisfaction operations that may be applied. If there is no bias then the order of operations should not matter, but the efficiency may vary greatly. Much of the previously reported work in PROTEAN went towards developing the BB1 framework for using heuristics to opportunistically determine the best operation to apply at any given time. The BB1 framework has been used to control the solid abstract level of problem solving, but it has not yet been integrated with the atomic level functionality, which is currently manually controlled.

Atomic level constraint satisfaction operations. As part of his Ph.D. work Bruce Duncan has extended many of the solid level constraint satisfaction operations to the atomic level. The computations make use of three different computational techniques: the systematic search of conformational space, a probabilistic updating technique, and a random search (Monte Carlo) technique.

The systematic search procedures use constraint satisfaction techniques and attempt to completely sample the protein conformational space. This is done by examining the positions of small groups of atoms to delete positions that do not satisfy the applied constraints. This is an extension of the techniques used by the solid level of the PROTEAN system.

Since there the primary limitation of the systematic search technique is that not enough atom positions can be simultaneously examined, the probabilistic technique attempts to infer the likely atom positions. In this technique, the atom positions are represented as sampled spatial probability distributions. Distance constraints are also modeled as probability distributions. An updating procedure modifies the spatial distributions of the atoms using the the constraint distribution. This procedure is related to the probabilistic methods developed by Russ Altman. Heuristic constraints on the structure were determined by analyzing known structures.

The Monte Carlo technique computes instances of the protein conformation that span the allowed range. Although probabilistic techniques can be used to compute structures represented by the mean atomic positions and the atomic covariance matrices, detailed analysis requires that individual instances are produced. At the atomic level, the technique of systematically searching the protein conformational space is ineffective to compute instances of the protein conformation. To overcome this limitation, techniques based on random sampling of the conformational space are being developed. This procedure makes use of several computational techniques used by the atomic level system such as simplified representations and the application of constraint satisfaction techniques. These techniques are effective because at this stage of the computation, the protein conformational space has been approximately determined. This technique is has the advantage that the

atomic positions are selected based on a random number generator thus ensuring that an unbiased set of conformations are computed.

This work has made use of the three SUMEX SUN 3/75 workstations (clienta, clientb, clientc) and the SUN file server (KNIFE). An interesting feature of the system is that it can make use of the three SUN workstations operating in parallel (although the interface is a bit clumsy). The code is written in the C language or in the C-shell UNIX command interpreter.

Probabilistic operations at the atomic level. Russ Altman, as part of his Ph.D. thesis, has been working to develop an alternative method for representing the accessible volumes of atoms. Currently, we represent accessible volumes as lists of xyz locations sampled on a regular grid at a specified resolution. In Russ's approach accessible volumes are represented as probability density functions described by a mean and covariance matrix. Constraints are also represented by a mean and variance, and a Kalman filter, which is basically another form of Bayes' formula, is used to determine the reduction in accessible volume. This approach has been used to reconstruct the amino acid tyrosine, and is being extended to handle larger numbers of atoms.

D. Relevant Publications

- 1) Altman, R. and Jardetzky, O.: *New strategies for the determination of macromolecular structures in solution.* Journal of Biochemistry (Tokyo), Vol. 100, No. 6, p. 1403-1423, 1986.
- 2) Altman, R. and Buchanan, B.G.: *Partial Compilation of Control Knowledge.* Proceedings of the AAI, pp 399-404, 1987.
- 3) Altman, R., Duncan, B., Brinkley, J., Buchanan, B., Jardetzky, O.: *Determination of the spatial distribution of protein structure using solution data,* Proceedings of the Alfred Benzon Symposium 26: NMR Spectroscopy and Drug Development, Copenhagen, 1987.
- 4) Brinkley, J., Cornelius, C., Altman, R., Hayes-Roth, B. Lichtarge, O., Duncan, B., Buchanan, B.G., Jardetzky, O.: *Application of Constraint Satisfaction Techniques to the Determination of Protein Tertiary Structure.* Report KSL-86-28, Department of Computer Science, 1986.
- 5) Brinkley, James F., Buchanan, Bruce G., Altman, Russ B., Duncan, Bruce S., Cornelius, Craig W.: *A Heuristic Refinement Method for Spatial Constraint Satisfaction Problems.* Report KSL 87-05, Department of Computer Science.
- 6) Brinkley, J.F., Altman, R.B., Duncan, B.S., Buchanan, B.G., and Jardetzky, O.: *The heuristic refinement method for the derivation of protein solution structures: Validation on cytochrome-b562,* KSL Technical Report 88-03, 1988.

- 7) Brugge, J.A., Buchanan, B.G., and Jardetzky, O.: *Toward automating the process of determining polypeptide secondary structure from ^1H NMR data*, To be published in J. Computational Chemistry, 1988.
- 8) Buchanan, B.G., Hayes-Roth, B., Lichtarge, O., Altman, A., Brinkley, J., Hewett, M., Cornelius, C., Duncan, B., Jardetzky, O.: *The Heuristic Refinement Method for Deriving Solution Structures of Proteins*. Report KSL-85-41. October 1985.
- 9) Duncan, B., Buchanan, B., Hayes-Roth, B., Lichtarge, O., Altman, R., Brinkley, J., Hewett, M., Cornelius, C., and Jardetzky, O.: *PROTEAN: A new method for deriving solution structures of proteins*, Bull. Mag. Res., 8:111-119, 1987.
- 10) Garvey, Alan, Cornelius, Craig, and Hayes-Roth, Barbara: *Computational Costs versus Benefits of Control Reasoning*. Report KSL 87-11, Department of Computer Science.
- 11) Hayes-Roth, B.: *The Blackboard Architecture: A General Framework for Problem Solving?* Report HPP-83-30, Department of Computer Science, Stanford University, 1983.
- 12) Hayes-Roth, B.: *BB1: An Environment for Building Blackboard Systems that Control, Explain, and Learn about their own Behavior*. Report HPP-84-16, Department of Computer Science, Stanford University, 1984.
- 13) Hayes-Roth, B.: *A Blackboard Architecture for Control*. Artificial Intelligence 26:251-321, 1985.
- 14) Hayes-Roth, B. and Hewett, M.: *Learning Control Heuristics in BB1*. Report HPP-85-2, Department of Computer Science, 1985.
- 15) Hayes-Roth, B., Buchanan, B.G., Lichtarge, O., Hewett, M., Altman, R., Brinkley, J., Cornelius, C., Duncan, B., and Jardetzky, O.: *PROTEAN: Deriving protein structure from constraints*. Proceedings of the AAAI, 1986, p. 904-909.
- 16) Jardetzky, O.: *A Method for the Definition of the Solution Structure of Proteins from NMR and Other Physical Measurements: The LAC-Repressor Headpiece*. Proceedings of the International Conference on the Frontiers of Biochemistry and Molecular Biology, Alma Alta, June 17-24, 1984, October, 1984.
- 17) Lichtarge, Olivier: *Structure determination of proteins in solution by NMR*. Ph.D. Thesis, Stanford University, November, 1986.
- 18) Lichtarge, Olivier, Cornelius, Craig W., Buchanan, Bruce G., Jardetzky, Oleg: *Validation of the First Step of the Heuristic Refinement Method for the Derivation of Solution Structures of Proteins from NMR Data.*, Proteins: Structure, Function and Genetics, 2:340-358, 1987.

E. Funding Support

Title: Interpretation of NMR Data from Proteins Using AI Methods

PI's: Oleg Jardetzky and Bruce G. Buchanan

Agency: National Science Foundation

Grant identification number: DMB-8402348

Total Award Period and Amount: 2/1/87 - 9/30/89 \$120,000 (includes direct and indirect costs)

Current award period and amount: 2/1/87 - 9/30/89 \$120,000 (includes direct and indirect costs)

III. RESEARCH PLANS

A. Goals & Plans

The AI research phase of the PROTEAN project is being terminated with the completion of thesis work by Bruce Duncan and Russ Altman. In the present stage of the PROTEAN project, most of the effort is being directed towards a successful application of the system in the Stanford Magnetic Resonance Laboratory of Prof. Oleg Jardetzky (SMRL). All of the PROTEAN system calculations will be performed on national supercomputing resources or on mini-supercomputers assembled in the SMRL for this purpose. Additional program development will be done on workstations, communicating via the local area network maintained by SUMEX.

Input data generation and output data processing also require a certain amount of additional computing, briefly described in the following.

1. SOLUTION STRUCTURE DETERMINATION

Input Generation

For simplicity of computation, hydrogen atoms which have no independent degrees of motional freedom are included in pseudo-atoms. "Pseudo-structures" of the twenty common amino acid residues are described by K. Wutrich, M. Billeter and W. Braun in J. MOL. BIOL. (1983) 169,949-961. The inclusion of the specific amino acid pseudo atoms list for a protein is automated.

Initial Mean Positions / Variance-Covariance Matrices

- (i) Initial means and variance-covariance values of the CA atoms are known from a previous Kalman filter update run.
- (ii) The mean positions of the added backbone, and amino acid residue pseudo atoms, are derived randomly from the CA values of (i). The initial variances of the mean positions are consistent with the uncertainty of the atom being anywhere within the protein domain. Zero covariances are assumed initially between any of two atom positions. Initial CA variances are taken from (i).

Generation of Distance Constraints

- (i) Distance Ranges. The set of distance constraints between pseudo atoms to be consistently updated by the Kalman filter mechanism consists of:

Constraints determined by the primary and secondary structure:

1. Covalent bonding within the backbone, within amino acid residues, and of those to the backbone.
2. Nonbonded distances in the backbone determined by covalent angles, e.g. for residue i:

$$\text{CA}(i) \quad \text{N}(i+1) = 2.38 \text{ \AA}$$

$$\text{CA}(i) \quad \text{O}(i) = 2.40 \text{ \AA}$$

$$\text{O}(i) \quad \text{N}(i+1) = 2.27 \text{ \AA}$$

3. Hydrogen bonding within helices.
4. Short range distances of protons in amino acid groups (pseudo-atoms) to the backbone protons which are three or less dihedral angles away.

Long range NMR NOE distance constraints.

- (ii) Variances. The distance constraints variances are evaluated as a sum of the individual pseudo-atoms variances, and the distance variance.

1. A pseudo-atom variance is assumed to follow the variance of a uniform distribution, i.e. (its van der Waals radius (e.g. 1.8 Å for a CH₃ group)/5)**2.
2. Similarly, if a distance range variance is assumed uniform, it is evaluated by ((a-b)**2)/12 where [a,b] is the distance range.

Thus, the variance of a long range NOE distance constraint is assumed to equal 1.1 Å (= 2* ((1.8)**2 / 5) + ((6 - 3)**2) / 12); the distance includes correction for protons, as discussed in the following).

- (iii) Distance Correction. Since NMR NOE distance constraints are observed between protons, short- and long-range NOE distance constraints have to be corrected for the C-H bond length of CA. 1Å.

In order to generate appropriate means and covariances, and the corresponding set of distance constraints for the Kalman filter program, some computational resources are to be utilized.

Output Data Processing

Results of a Kalman filter run have to be evaluated in terms of the structure's consistency with the distance constraints, and with structures obtained by other methods. This evaluation entails some computations, e.g. distance calculations given a set of (x,y,z) coordinates.

BLOCH EQUATION CALCULATIONS

Input Generation

For the input to the bloch equation program, a standard pdb file is utilized for input data generation :

- (a) A data set including only protons is generated.
- (b) Methyl groups are identified and average positions are evaluated.
- (c) A distance matrix is calculated, taking into account only those proton pairs which are less than 6 Å apart.
- (d) The number of all possible NOE distances for each proton is calculated.

Output Evaluation

A large amount of data processing on an output file has to be performed, as is shown in the following example.

Given the NOE output file for the protonated trp repressor, which was recently calculated, the following computations were performed:

- (a) NOE intensities of proton pairs for various mixing times are separated.
- (b) For each mixing time the following NOE's were to be identified:
 - (1) Intensities of HN(i)-HN(i+1) for each i, where i is the residue number.
 - (2) Intensities of each HN(i) other than to other HN's or HO's.
 - (3) Intensities of (2), but also other than those to side chain i.
 - (4) Intensities of (3), but also other than sequential.

B. Justification for Continued SUMEX Use

The communication and other capabilities provided by SUMEX and the SUMEX staff is essential for continued success of such a large and evolving distributed project.

C. Other computing resources

At this time our computational resources are almost adequate as long as we can continue to use the SUMEX-supported local area network. Since PROTEAN is evolving into a distributed problem-solving system, we actively support and require extensive network access capabilities. We believe that the requirements of PROTEAN make it an ideal test problem for many of the distributed systems issues being considered by the SUMEX core research staff.

IV.A.6. Reasoning Under Uncertainty

Pragmatic Approaches to Reasoning under Uncertainty in Medical Decision Support Systems

Principle investigator: E.H. Shortliffe

Project manager: G.F. Cooper

Students: R.M. Chavez, D.E. Heckerman, E.H. Herskovits,

E.J. Horvitz, M. Shwe, and H.J. Suermondt

I. Summary of research program

I.A. Overview

Medicine is replete with uncertainty. In particular, there is uncertainty due to incomplete and inexact scientific models of human health and disease, and there is uncertainty secondary to incomplete and erroneous data about individual patients. Therefore, in general, it is important that computer systems that assist in medical decision making be capable of representing and reasoning with uncertainty. We are exploring the use of probability theory as a representation of uncertainty in medical diagnostic systems. The advantages of using a probabilistic representation, include that probability 1) is mathematically well-defined and has been studied extensively, 2) provides a common, well-established language for communicating uncertainty, 3) allows the combination of subjective probabilities from medical experts with statistics gathered from databases, and 4) is a necessary component of normative decision systems that are based on decision theory. Nonetheless, there are potential problems associated with using a probabilistic representation. Key challenges include developing knowledge acquisition methods and probabilistic inference techniques that are tractable, particularly in light of the large amount of knowledge that must be represented in broad areas of medicine.

We are currently pursuing approaches that we believe hold promise in making practical the probabilistic reasoning in medical expert system. The specific probability representation we are exploring is called a Bayesian belief network. Such networks provide a mechanism for intuitively and efficiently specifying the probabilistic dependencies among diseases, intermediate pathophysiological states, and evidence (e.g., symptoms). In Section I.B we discuss three areas of current belief network research. Section I.B.1 summarizes recent work on methods that enable the computer to control its own probabilistic reasoning processes so that it focuses its potentially scarce computing resources (particularly time) on the most critical diagnostic inference tasks. Section I.B.2 highlights some specific probabilistic inference algorithms that we are developing and evaluating. Here again the critical issue is speed of probabilistic inference. Section I.B.3 discusses techniques we have developed for efficient acquisition of probabilistic knowledge from medical experts.

I.B. Research progress

I.B.1. Reasoning about inference tradeoffs

Recent work on inference tradeoffs has focused on the analysis of techniques for balancing the costs and benefits of decision-theoretic inference. In this work, we use decision-theoretic control of decision-theoretic inference as a model of rational computational decision making. In particular, the computer system continually calculates the expected value of additional computation in deciding whether to make a diagnostic suggestion based on current computation or to continue computing [16]. For example, the amount of time taken by the computer to refine its diagnosis should generally be small when the patient has an obviously life-threatening problem. Nonetheless, the computer must compute long enough to attain a reasonably accurate diagnostic suggestion, otherwise, based on the system's suggestion, the patient might be treated for the wrong underlying condition. Developing computer systems that can dynamically make tradeoffs between computation time and diagnostic accuracy is a difficult problem. We have made use of partial characterizations of decision-theoretic inference to reason about the value of continuing to compute versus making a diagnostic suggestion. We particularly have been studying the value of applying a probabilistic bounding algorithm to diagnostic problems. This work makes use of summary information about the convergence of interval bounds on the probabilities of potential diseases. Related work has focused on using decision-theory to control the construction of a relatively small, context-sensitive decision model from a much larger belief network, so that the small model can be efficiently solved. The QMR-DT knowledge base (see Section I.B.2) is being used as the large belief network.

I.B.2. Efficient probability inference algorithms

During the past year, we have continued to explore algorithms for probabilistic inference using Bayesian belief networks. In particular, we have implemented inference algorithms that are exact, approximate, and special-case. Exact algorithms return precise, accurate probability answers, but these algorithms can be slow when applied to large, complex medical domains. Approximation algorithms tradeoff precision and accuracy for speed, but bound the imprecision and the inaccuracy with explicit error terms. Special-case algorithms take advantage of special properties of some classes of belief networks (such as restrictions on topology) in order to solve inference problems efficiently. During the past year we have begun to evaluate these exact, approximate, and special-case algorithms.

Exact Algorithms

Regarding exact probabilistic inference methods, we have implemented the Lauritzen-Spiegelhalter algorithm. In the process, we have gained insight into the graphical manipulations necessary for efficient belief-network