Research in Progress

Artificial Intelligence Research at the University of Maryland

Jack Minker

Larry S. Davis

Computer Science Department
University of Maryland
College Park, MD 20742

THE UNIVERSITY OF MARYLAND'S Computer Science Department conducts a broad research program in both theoretical and applied artificial intelligence. Nine faculty and more than fifty research associates and graduate students are involved in AI research. Projects are funded by a large number of government agencies, as well as by several major corporations. The computing environment will improve dramatically over the next several years, due in large part to a Coordinated Experimental Research Equipment Grant awarded to the Computer Science Department by the National Science Foundation in 1982.

In addition to the research program in AI, the Department offers a large number of courses at both the graduate and undergraduate levels on all facets of AI. The principal AI laboratories also sponsor numerous colloquia by visiting scientists and permanent laboratory personnel. The principal research areas are computer vision, search and decision making, parallel problem solving, and database research.

Research in computer vision spans a broad spectrum of fundamental and applied problems. The basic research in vision is in long-term support of applications in industrial automation, autonomous vehicle navigation, smart reconnaissance and remote sensing of the environment. A multi-faceted approach to vision research is supported in a community including scientists trained in computer science, mathematics, physics and psychology, and electrical engineering.

Research in expert systems and intelligent search is conducted in the Machine Intelligence and Pattern Analysis Laboratory. The principal research areas include parallel implementation of hierarchical search procedures, the theory of search for game trees and the design of expert systems based on a "set-covering" paradigm rather than a production rule paradigm.

Finally, in the area of problem solving and database research, a substantial effort is being made to design and build parallel problem solving systems capable of analyzing problems whose solution requires not only deep inference, but also access to large databases of information. A first generation parallel problem solving system, called PRISM, has been developed and is currently being implemented on a multimicroprocessor computer system. Fundamental problems in automatic inference, involving non-monotonicity and circumscription, are also being studied in the AI and Database laboratory.

Computer Vision

Research in computer vision is conducted in the Computer Vision Laboratory [CVL], which is part of the University's new Center for Automation Research [CFAR]. Azriel Rosenfeld is the Director of CFAR, and Larry Davis is the head of CVL. The laboratory supports over 30 graduate students from the Computer Science Department, and has its

own computing laboratory that includes a VAX 11/780 and specialized equipment for image input, output and display.

Research in CVL spans a broad range of fundamental problems and applications. The principal research projects in the laboratory are:

1. *Multiresolution image representations:* Azriel Rosenfeld has been studying the roles that multiresolution image representations, such as pyramids, play in computer vision systems both at lower levels of processing where highly structured image representations need to be developed, and at higher levels where a continuous interface is required between pixel and region level image models and the symbolic models ordinarily used to represent scene structure. Many of these issues were addressed in a recent doctoral dissertation by Hong (1982).

2. *Shape from shading, texture, shape·* Chia-Hoang Lee (1983) has recently completed a doctoral thesis on shape from shading; Lee stressed the importance of considering the effects of image sampling and digitization on shape from shading models, and developed numerically more stable algorithms for determining the shape of matte objects Stanley Dunn is currently working on a doctoral thesis on shape from shape and texture. His early work (Davis & Dunn 1983) was based on the model developed by Witkin at M.I.T ; he developed several efficient algorithms for computing maximum likelihood estimates of planar surface orientation from observed image texture and has recently completed an empirical investigation of shape from texture using natural imagery

3. *Expert Vision:* Takashi Matsuyama (visiting scholar from Kyoto University) has been working with Vincent Huang on an expert system for understanding aerial images; they are extending the architecture developed by Matsuyama in his thesis to allow for feedback from symbolic analysis to the low level vision operators, and to support more sophisticated representations of image structure (Hwang, Matsayuma, Rosenfeld, & Davis 1983) Larry Davis is studying the problem of matching images taken under very dissimilar conditions This project is based on matching intermediate level image representations and utilizes multiresolution representations of salient image structures (Davis, Kitchen, Hu, & Hwang 1983)

4. *Time-varying imagery:* Allen Waxman has completed a thorough analysis of the recovery of structure and motion from information contained in optic flow fields (Waxman & Ullman 1983) He is now working with Kwangyoen Wohn on developing the theory and algorithms to construct the image motion representation required to compute structure and object motion from time-varying images. Larry Davis and Allen Waxman are working with Nader Kazor on an expert system for target tracking that includes both static and dynamic models of the targets and their environment, and attempts to incrementally construct a model of the scene based on the interactions of static and dynamic objects.

5. *Object recognition:* Larry Davis and Gyorgy Fekete are working on an approach to recognizing 3-D objects from two-dimensional images based on a property sphere representation for the objects (Fekete & Davis in preparation) Property spheres encode absolute geometric and topological properties of the appearance of a shape as a function of viewpoint. Larry is also working with Terry Silberberg on another approach to three-dimensional object recognition that can recognize objects even when they are only partially visible in the image (Silberberg, Davis, & Harwood 1983)

6. *Parallel processing:* Azriel Rosenfeld has been investigating parallel algorithms for image analysis for over 15 years In collaboration with Angela Wu (of American University) and I V Ramakrishnan, he is currently studying general graph models for parallel computation, with applications to image analysis and processing

7. *Low-level vision:* Low-level vision has been a principle research topic in the lab for over twenty years David Harwood and Stanley Dunn (Dunn, Harwood, & Davis 1983) developed a theory of "uniform-error" thresholding and applied it to a wide variety of images. David Harwood and Muralidhara Subbarao (Harwood, Subbarao, & Davis 1983) introduced rank-order correlation masks for image texture analysis and showed that very high image recognition rates could be achieved with their approach

8. *Data structures for vision* Hanan Samet directs a group that is concerned with hierarchical data structures for image analysis and cartography. They have done extensive work on quadtree-like structures for storing and manipulating point, line and region data (Samet 1983)

9. *Autonomous vehicle navigation:* CVL, in collaboration with the Westinghouse corporation, is developing a testbed facility for autonomous ground vehicle navigation, with emphasis on visual control. The vehicle will integrate analyses of data from a variety of sensors (including direct ranging and optical video) to navigate a vehicle over natural, rugged terrain

References

Davis, L S , Janos L., & Dunn, S (1983) Efficient recovery of shape from texture. *IEEET-Pattern Analysis and Machine Intelligence*, 5, 485-492.

Davis, L S., Kitchen, L., Hu, F , & Hwang, V.October (1983) Image matching using generalized Hough transforms Center for Automation Research TR-27.

Dunn S., Harwood D., & Davis, L S (1983) Local estimation of the uniform error threshold. University of Maryland Computer Science TR-1279.

Fekete, G., and Davis, L S. (in preparation) Property spheres: A new representation for 3-D object recognition. Center for Automation Research.

Harwood D., Subbarao, M , & Davis, L. S (1983) Texture classification by local rank correlation. University of Maryland Computer Science TR-1314.

Hong, T H. (1982) *Multiresolution image analysis* Doctoral dissertation, Department of Computer Science, University of Maryland.

Hwang, V , Matsayuma, T , Rosenfeld, A., & Davis, L (1983) Evidence accumulation for spatial reasoning in aerial image understanding. Center for Automation Research TR-28.

Lee. C. H (1983) *The use of shading in image segmentation and surface orientation estimation.* Doctoral dissertation, Department of Computer Science, University of Maryland.

Samet H (1983) The quadtree and related hierarchical data structures. University of Maryland Computer Science TR-1329

Silberberg T., Davis, L. S , & Harwood, D. (1983) An iterative Hough procedure for three dimensional object recognition. Center for Automation Research TR-20

Waxman, A , and Ullman, S (1983) Surface structure and 3-D motion from image flow. Center for Automation Research TR-24.

Search and Decision Making

The Machine Intelligence & Pattern Analysis Laboratory [MIPAL] is the focus for research on search and decision processes. The Laboratory is directed by Laveen Kanal; other faculty associated with MIPAL are Dana Nau and James Reggia, and Visiting Assistant Professor G.R. Dattatreya. The principal research projects in MIPAL are:

1. *Search:* Laveen Kanal and his students have developed non-directional (top-down, bottom-up, non-left-right) search algorithms for searching attributed AND/OR graph models and have applied them to signal understanding (Stockman & Kanal 1983) and image understanding (Kanal, Lambird, & Lavine 1983) A general branch and bound model has been developed and shown to subsume many search algorithms for searching state space and AND/OR graph models (Kumar & Kanal 1983) Adaptation of such algorithms for implementation under parallel hardware is also under investigation (Kanal & Kumar 1981). Recently, Kumar and Kanal (1983) introduced the concept of a "composite" decision process which unifies much of the previous disparate theory concerning search procedures and has led to new insights concerning the relationship among branch and bound, dynamic programming and heuristic search procedures

2. *Hierarchical decision procedures:* Kanal has been studying the design and optimization of hierarchical statistical classifiers for many years (Kanal 1977). Recent work in the Laboratory on this topic includes methodology for adaptively improving the performance of pattern recognition decision trees (Dattatreya and Kanal 1983)

3. *Decision making on game trees:* Until recently, it was believed that the deeper a game tree was searched, the better the decision would be However, Dana Nau showed, several years ago, that this is not always the case, but that deeper search is sometimes worse than shallow search (Nau 1980). He has recently extended this research by identifying various factors governing whether or not this pathological situation will occur (Nau 1983), and is exploring alternatives to minimaxing which can, under some circumstances, avoid such pathological behavior

4. *Automated manufacturing:* Nau is currently investigating the possible use of A I techniques in automating a manufacturing task known as process selection (Nau & Chang in press) He is also investigating the kinds of spatial reasoning that are required for process planning.

5. *Expert systems:* James Reggia and Dana Nau have been exploring an approach to problem solving based on a generalization of the set covering problem (Reggia, Nau, & Wang in press) A reasoning system based on this approach has been implemented and successfully used to create expert systems for several problem domains. James Reggia has developed a prototype, domain independent knowledge engineering tool called KMS (Reggia & Perricone 1982); it is a high-level, non-procedural knowledge representation language designed for use by application specialists. The system is being used by computer-inexperienced medical students and physicians at the University of Maryland Medical School. Reggia has also recently completed an answer justification system that can "explain" the posterior probability distribution assigned to outcomes by a decision aid based on Bayesian classification (Reggia & Perricone 1983)

6. *Error modelling:* James Reggia, jointly with Dana Nau and faculty from the medical school, is currently studying the simulation and classification of errors occurring in discrete sequential processes A heuristically-guided prototype simulator has been implemented and is currently being evaluated (Reggia & Ahuja 1983).

References

Dattatreya, G. & Kanal, L. (1983) *Adaptive pattern classification with random costs and its applications.* University of Maryland Computer Science TR-1348, November.

Kanal, L. (1977) Hierarchical classifier theory and interactive design In P. Krishnaiah (Ed.), *Application of Statistics* North-Holland.

Kanal, L. & Kumar, V. (1981) Parallel implementations of a structural analysis algorithm. *Pattern Recognition and Image Processing,* 452-458, August.

Kanal, L., Lambird, B , & Lavine, D (1983) Expert systems for understanding remotely sensed images. Proceedings, the Army Conference on Application of Artificial Intelligence to Battlefield Information Management, Naval Surface Weapons Center, Washington D.C , April

Kumar, V. & Kanal, L. (1983a) A branch and bound approach to understanding and synthesizing AND/OR graph search. *Artificial Intelligence,* 21, 179-198.

Kumar, V. & Kanal. L. (1983b) The composite decision process: A unfying formulation for heuristic search, dynamic programming and branch and bound procedures. *AAAI-83*, Washington D.C., 220-224, August.

Nau, D. S (1980) Pathology of game trees: A summary of results. *AAAI-80*, 102-104.

Nau, D. S. (1983) Pathology on game trees revisited, and an alternative to minimaxing. *Artificial Intelligence*, 21, 221-224.

Nau, D S. & Chang, T. C. (in press) Prospects for process selection using artificial intelligence. *Computers in Industry*.

Reggia, J., Nau, D. S., & Wang. P. (in press) Diagnostic expert systems based on a set covering model. *International Journal of Man-Machine Studies*.

Reggia, J , & Perricone. B. (1982) *KMS reference manual.* University of Maryland Computer Science TR-1136

Reggia, J & Ahuja, S. (under review) Simulation of phonemic errors using artificial intelligence symbol processing techniques

Reggia, J. & Perricone, B. (under review) Answer justification in medical decision support systems based on Bayesian classification.

Stockman, G., & Kanal, L. (1983) Problem reduction representation for the linguistic analysis of waveforms. *IEEET-Pattern Analysis and Machine Intelligence*, 5, 287-298.

Problem solving and Database Research

The AI and Database Research Laboratory is directed by Professor Jack Minker. Other faculty connected with the lab are Donald Perlis and John Grant. Perlis is concerned with theoretical issues associated with logic programming, non-monotonic reasoning, and cognitive problems. Grant, from Towson State University, is primarily concerned with theoretical issues in databases. The laboratory's principal research projects are:

1. PRISM: PRISM is a parallel inference system which will run on the ZMOB computer. ZMOB was designed by Chuck Rieger, and consists of 256 Z80A microprocessors interconnected on a high-speed conveyor belt, with the VAX 11/780 as the host processor. Moblets (individual Z80A's) can be accessed either by physical address or by pattern matching ZMOB supports point-to-point high speed transmission between processors. PRISM has been implemented and tested on a simulated system, and is now being transported to the ZMOB hardware.

   The underlying ideas behind PRISM appeared in (Eisinger, Kasif, & Minker 1982), and a detailed functional description of the system appears in (Minker, Asper, Cao, Chakravarthy, Chesapeake, Kasif, Kohli, & Piazza 1982). PRISM consists of several parts: a user host interface that exists on the VAX; a set of Z80A machines designated as problem solvers; a set of machines designated as Extensional Database Machines that store ground atomic formulas (relational database tables); a set of machines designated as Intensional Database Machines that store procedures (the general rules in the system); and an Intensional Database Monitor machine which

determines if the Intensional Database machines are overloaded.

The user specifies control in terms of the sequence of atoms to be executed in a set of problems to be solved. Atoms (and, similarly, procedures) can be executed in parallel, sequentially, or as specified by a partial ordering. The user may specify the configuration (i.e , the number of moblets required at a minimum) in which a problem is to be run. If additional moblets are available, PRISM is designed to automatically integrate them into its problem solving activities.

2. *Control structures:* Controlling searches in a logic program is crucial for developing efficient expert systems and problem solvers. Kohli and Minker (1983) discuss how failures to find solutions along paths, together with integrity constraints, may be used to provide more intelligent search than is achievable with backtracking The control structure for PRISM is described in (Kasif, Kohli, & Minker 1983). A major objective of the research in this area is the investigation of control issues in distributed problem solving. Preliminary studies have been directed towards developing a control specification language for logic programming systems.

3. *Interfacing between predicate logic programs and large databases:* AI programs will have to deal effectively with large databases. In Chakravarthy, Minker, and Tran (1982) alternative methods of interfacing a predicate logic program with a relational database in a system (*e.g.,* PROLOG) are described. Approaches discussed are modifications to PROLOG interpreters and meta-level constructs. In Grant and Minker (1983) the optimization of compiled deductive searches that do not include recursion is considered. A branch-and-bound type algorithm is used that takes advantage of a database that is indexed and the number of responses anticipated. It performs the union of a set of conjuncts in an optimal fashion.

4. *Generalized closed world assumption:* The closed world assumption, identified by Reiter and discussed by Nicolas and Clarke for Horn clauses, was generalized in Minker (1982). A proof theoretic and a model theoretic definition were provided, and it was shown that the definitions were equivalent. The generalization leads to a non-monotonic logic

5 *Recursive algorithms:* When recursive algorithms are part of a deductive database, problems may arise with respect to termination. Minker and Nicolas (1983) show how one can obtain termination conditions for classes of clauses that contain recursive axioms.

6. *Database theory:* The notion of numerical dependencies is introduced in Grant and Minker (1983) It formalizes the constraint that with an element of an attribute, or set of attributes, at most k elements of another attribute can be associated. The uses of numerical dependencies in database design are considered via generalized normal forms.

7. *Non-monotonic logic:* Minker and Perlis are also investigating the semantics of circumscription in two aspects. First, they extended McCarthy's notion of circumscription so that certain desired individuals may be protected by circumscription, thus allowing indefinite data. In Minker and Perlis (1983) models are defined for protected circumscription that are generalizations of McCarthy's models, and it is shown that McCarthy's soundness theorem holds for these new models as well. Secondly, they have sought completeness theorems guaranteeing that sentences true in all models (e.g., McCarthy's or their protected one) are also circumscriptively provable (protected or otherwise). Such results have been found for the case of axioms with finite Herbrand bases (no function letters). Furthermore, their proofs precisely delineate which cases of the circumscriptive schema are needed. They are currently involved in extending these results to certain special cases of infinite Herbrand bases Minker and Perlis (1983) have applied their results to problems in databases.

8. *Truth, logic and cognitive modeling:* The need to represent process information leads to a kind of self-reference in which a reasoning system must reason about its own behavior. This can lead to contradictory axioms, *a la* Russell's paradox or the Liar paradox An appropriate representation of quotation and unquotation (or truth) has been developed by Don Perlis that avoids contradiction, yet is natural, quite broad in expressive power, and computationally tractable (Perlis under revision). Applications to AI and databases are considered in Minker and Perlis (1983). A pilot computer model is being developed, based on these ideas, to explore real-time common sense reasoning in the context of short and long term memory.

References

Chakravarthy, U., Minker, J , & Tran, D. (1982) Interfacing predicate logic languages and relational databases. *Proceedings of the First International Logic Programming Conference,* Marseilles, September.

Kasif, S., Kohli, M. & Minker, J. (1983) Control structure of PRISM: A parallel inference system for problem solving. In Proceedings of the Logic Programming Workshop, Faro Portugal, 123-152, June.

Kohli, M. & Minker, J. (1983) Intelligent control using integrity constraints. *AAAI-83,* 202–205.

Minker, J., Cao, D , Chakravarthy, U , Csoeke-Poeckh, A., Kasif, S., Kohli, M., Piazza, R., & Wang, D. Parallel problem solving on ZMOB. (1983) *Proceedings, Trends and Applications 83,* 142-146.

Grant, J. & Minker, J (in press) On optimizing the evaluation of a set of expressions *International Journal of Computer and Information Systems.*

Minker, J. (1982) On indefinite databases and the closed world assumption. *Proceedings, Conference on Automated Deduction,* 6.

Minker, J & Nicolas, J. M. (1983) On recursive axioms in deductive databases. *Information Systems,* 7, 4.

Minker, J & Perlis, D. (1983) On the semantics of circumscription. *University of Maryland Computer Science TR-1341*

Minker, J & Perlis, D (1983) Applications of protected circumscription. *University of Maryland Computer Science TR-1342,* September

Perlis, D. (under revision) A foundation for knowledge representation *Artificial Intelligence.*

Eisinger, N., Kasif, S. & Minker, J. (1982) Logic Programming: A Parallel Approach. *Proceedings of the First International Logic Programming Conference,* Marseilles, September

Minker, J., Asper, C., Cao, D., Chakravarthy, U S., Chesapeake, A , Kasif, S., Kohli, M , & Piazza, R. (1982) Functional Specification of the ZMOB Parallel Problem Solving System. *University of Maryland Computer Science Technical Note Z-1,* August.