
RESEARCH IN PROGRESS

Artificial Intelligence Research at the University of California, Los Angeles

Michael G. Dyer, Editor

Computer Science Department, University of California, Los Angeles, California 90024

Research in AI within the Computer Science Department at the University of California, Los Angeles is loosely composed of three interacting and cooperating groups: (1) the Artificial Intelligence Laboratory, at 3677 Boelter Hall, which is concerned mainly with natural language processing and cognitive modelling, (2) the Cognitive Systems Laboratory, at 4731 Boelter Hall, which studies the nature of search, logic programming, heuristics, and formal methods, and (3) the Robotics and Vision Laboratory, at 3532 Boelter Hall, where research concentrates on robot control in manufacturing, pattern recognition, and expert systems for real-time processing.

Current Research Projects at the UCLA Artificial Intelligence Laboratory

In early 1984, the UCLA Computer Science Department was awarded a private foundation grant to establish of a research laboratory in artificial intelligence. In addition, the School of Engineering and Applied Sciences agreed to provide matching funds and floor space to house the new laboratory. The lab became operational in July, 1984 and fifteen Apollo workstations (ten DN320s, four DN460s, and one color DN660) were purchased for the lab. Another six Apollos were purchased under a Computer Science Department NSF/CER grant. The AI lab is also funded by grants from the Hughes Artificial Intelligence Center.

Current research at the AI lab is concentrated in the area of natural language processing and cognitive modeling. In general, we are interested in gaining computational insight into the relationship of thought to language; the organization and content of human memory; how humans reason, argue, and persuade; representation of knowledge; and the processes of learning, imagination, and creativity. Our approach involves designing and building pro-

cess models which exhibit natural language I/O behavior similar to that produced by humans. Example texts or protocols are selected that are considered prototypic of a specific domain and that give rise to the basic cognitive issues being addressed.

The research program is divided up into a number of research projects, each consisting of the design and implementation of one or more computer programs. The projects are directed by Professors Dyer and Flowers and led by an advanced graduate student or research associate.

The MUSE Project

MINSTREL: Storytelling. MINSTREL is designed to tell stories in the domain of King Aurther and his knights of the round table. In MINSTREL, high-level storytelling goals work to make plan-based stories more interesting and purposeful. For instance, one heuristic for increasing the tension in a story is to make the main goal more important to the main character. Another heuristic is to delay achievement of an important preservation goal.

Key issues being explored are: (1) a theory of themes: their representation, selection and application; (2) the interaction between dramatic goals, episodic memory, and planning; and (3) the adaptation of memory episodes from one domain in service of dramatic goals for the storytelling domain.

Participants: Scott Turner (contact), Michael Dyer.

References: M. Dyer (1981), S. Turner & M. Dyer (1985).

DAYDREAMER: Imagination and Relaxed Hypothetical Planning. DAYDREAMER is a computer program currently under construction, that implements a computational theory of daydreaming. DAYDREAMER consists of:

- An imaginative scenario generation/planning mechanism.
- A personality model which provides the goals, some wishful and others realistic, that guide the scenario generation.
- An affect component, in which emotional states both initiate, and are influenced by, daydreams.
- A dynamic episodic memory of experiences.
- Knowledge of the domain of interpersonal relations and common everyday occurrences, including planning knowledge.

A prototype version of DAYDREAMER which produces several daydreams based on several experiences is currently running. The output consists of English text supplemented by program trace information and a graphics display of working memory. At present, input experiences are represented and incorporated into dynamic episodic memory by hand. In the future we will hook DAYDREAMER up to a story understander to enable it to gain new experiences by reading stories. Finally, we will explore the relationship of daydreaming to more deliberate thought processes such as conversation, storytelling, and invention.

Participants: Michael Dyer, Erik Mueller (contact).

References: M. Dyer (1983b), E. Mueller & M. Dyer (1985b).

SOB-SISTER: Writing Human-Interest News Stories. This project investigates how natural language is generated by modeling the cognitive processes used by journalists to write news stories. By modeling generation as the integrated processes of understanding and evaluating story contents as a journalist would, and considering the ways in which the story can be approached and presented (its "angle"), the problem of text generation becomes the problem of how interesting aspects of situations are recognized, and how this recognition can be used to organize generational information. One type of story angle is situational irony, where character and reader expectations are violated. Given a representational framework for the concept of irony, certain classes of irony, such as "just dessert" and "role reversal" ironies, can be identified. In SOB-SISTER, the recognition of irony is used to organize the events of a situation for its telling as an interesting story.

Participants: John Reeves (contact), Michael Dyer.

References: J. Reeves & M. Dyer (1985).

The MENTOR Project

AQUA: Modeling Advice Giving. This research examines advice giving: The process of understanding a problem description and suggesting potential solutions. AQUA is a computer program that attempts to understand natural language descriptions of various problems encountered by

novice UNIX users, combining its own knowledge of UNIX with its understanding of the user's problem description in order to explain the user's problem and to suggest a solution.

A user problem can be understood in terms of the user's goals, the plans the user attempted in order to satisfy these goals, and the states resulting from the use of these plans. From a user problem description, AQUA suggests a solution by comparing its plan for achieving the user's goal with the plan actually attempted by the user, informing the user of any relevant differences. Typically, these differences include such common user mistakes as selecting an incorrect plan, or forgetting an important step in a correct plan. When AQUA does not have a plan for a user's goal, it attempts to form a reasonable plan by using heuristics that allow the transformation of the user's goal into goals with known plans

Participants: Alex Quilici (contact), Michael Dyer.

SHERLOCK: An Inferencing Game. SHERLOCK is an intelligent tutoring game designed to explicate and teach strategies for successful inferencing. This game presents the user with a screen containing a goal box (an icon containing a concept or idea to be supported) and playing pieces (icons representing facts, beliefs, domain-specific rules, and general inference rules). Using a mouse, the player is able to move an icon to another part of the playing field, create a link between two icons, break a link, or ask for a hint. When the necessary antecedents for a particular icon are linked to it, an icon representing the consequent is automatically generated. In this way the player tries to build a graph which will support the goal node.

Because an inferencer is used, rather than a pattern matcher, the player is not limited to the "correct" answers anticipated by the author. Any valid inference chain will be recognized by the program. In addition this allows the program to make its own inferences and thus provide a wide range of feedback to the learner. Future plans include: Adding an intelligent coach to train less successful students in the use of these thinking strategies, and adding a playing-field editor to the game to create a universal intelligent tutoring system.

Participants: Rich Feifer (contact), Michael Dyer, Margot Flowers.

References: R. Feifer, M. Dyer, & M. Flowers (1985).

The LACQ Project

RINA: A Self-Extending Phrasal Lexicon. RINA is a computer program under construction, that models a second language speaker who learns new phrases from their context. RINA operates by accessing and modifying a phrasal lexicon of pattern/concept pairs. The program forms hypotheses from context about both the syntactic structure

and meaning of novel phrases. RINA stores generalizations of the episodic context in order to make available connotations for use in generating phrases that have been recently acquired.

Participants: Uri Zernik (contact), Michael Dyer.

References: M. Dyer & U. Zernik (1985), U. Zernik & M. Dyer (1985).

CHIE: Second Language Generation. This project investigates the ways in which knowledge of two languages is organized in the mind of a second language learner and how this knowledge is used in generating second language utterances.

CHIE is a program under development to model the generation of spoken English by a native speaker of Japanese. The program is designed to take input in the form of Japanese narrative text, and translate it into English text containing errors similar to those made by beginning Japanese learners. Work is focusing on: (1) the degree to which the two linguistic systems share knowledge; (2) how the speaker copes with gaps in lexical and grammatical knowledge; and (3) how formally learned knowledge of the second language differs from knowledge acquired in a natural way.

Participants: Michael Gasser (contact), Michael Dyer.

Reference: M. Gasser (1985).

The FLE Project: Foundations of Legal Expertise

In the field of law the decisions in previous cases often play a significant role in the presentation and outcome of new cases. Lawyers are constantly looking up old cases to aid them in preparing their own briefs. The key issues in this project, Foundations of Legal Expertise and Precedence-Based Legal Reasoning, are the representation of legal concepts and cases and the organization of a dynamic memory containing many instances of such concepts, cases, and general world knowledge.

We restrict our research to the domain of contract law and concentrate our initial effort on simple cases such as those presented to beginning law students. The first phase of this project is now underway. Its goal is the representation of a reasonable set of simple contract law cases and a dynamic memory which is able to receive new cases, integrate them into memory with existing cases, and discover related cases (precedents).

Participants: Seth Goldman (contact), Michael Dyer, Margot Flowers.

References: M. Dyer & M. Flowers (1985), S. Goldman, M. Dyer, & M. Flowers (1985).

The MORRIS Project: Moral and Reminding Inference Systems

CRAM: Causal Reasoning in Associative Memory. This

research addresses the issue of how people understand stories which contain planning failures. Planning failures are the result of poor planning or counterplanning on the part of one of the characters in a story. A program that can understand planning errors is important for two reasons: First, knowledge about bad planning can serve to help critique plans which are generated by a planner; and second, in domains where multiple agents are planning and counterplanning against one another, bad planning on the part of one character is often good planning on the part of another. For example, being deceived can be viewed either as a failure for one character or a success for another.

The knowledge structures used for recognizing planning errors are TAUs—Thematic Abstraction Units (Dyer, 1983). TAUs characterize the planning knowledge often expressed in adages such as:

A stitch in time saves nine.

Don't count your chickens before they're hatched.

Pride goeth before a fall.

Participants: Charles Dolan (contact), Michael Dyer.

References: C. Dolan (1985), C. Dolan & M. Dyer (1985a), M. Dyer (1983a), M. Dyer (1983c).

OCCAM: Explanation-Based Generalizations in Memory. A model of memory and learning is presented that indexes a new event by those features which are relevant to the explanation of why the event occurred. As events are added to memory, generalizations are created that describe and explain similarities and differences between events. The memory is organized so that when an event is added, events with similar features are noticed. An explanation process attempts to explain similar features and, if an explanation is found, a generalized event is created to organize the similar events. The explanation is then stored with the generalized event.

Participants: Michael Pazzani (contact), Michael Dyer, Margot Flowers.

Reference: M. Pazzani (1985).

The OpEd Project: Comprehension of Editorial Text

OpEd: Modeling Reasoning Comprehension. This project focuses on the knowledge structures and processes required for computer comprehension of opinions, arguments, and issues which arise in politico-economic editorials. The computational issues addressed here include:

- Representing general politico-economic knowledge.
- Recognizing beliefs and belief relationships.
- Following belief justifications.
- Representing beliefs and belief justifications.
- Dealing with interactions between reasoning and world knowledge.

Theoretical constructs and processing strategies for reasoning comprehension are currently being implemented in OpEd (Opinions to/from the Editor), a computer program that reads and answers questions about short politico-economic editorials. The design of OpEd is based on the demon-based conceptual parser and question answering model implemented in BORIS (Dyer, 1983), and argument graph developed by Flowers et al. (1982)

Participants: Sergio Alvarado (contact), Michael Dyer, Margot Flowers.

References: S. Alvarado (1985), M. Dyer (1983a), M. Dyer & W. Lehnert (1982), M. Flowers (1982), M. Flowers, M. McGuire & L. Birnbaum (1982).

JULIP: Understanding Analogies in Editorials. JULIP is a computer program which deals with the task of understanding analogies in the domain of editorial letters. Our theory of analogy comprehension draws upon knowledge of editorial comprehension, argumentation, and integrated natural language systems. JULIP currently works on a hand-coded conceptual representation of an editorial letter containing an analogy, and demonstrates its understanding of the analogy through question answering. In the future, the system should also be capable of generating analogies on its own, so that given one event, it can be reminded of past events that represent similar abstract concepts, even if the surface features of the events suggest that they represent different topics.

Participants: Stephanie August (contact), Michael Dyer, Margot Flowers

References: S. August & M. Dyer (1985), M. Flowers & M. Dyer (1984).

The EDISON Project

The goal of the Engineering Design Invention System Operating Naïvely (EDISON) project is the design of programs capable of creating novel mechanical devices. We start with simple mechanical objects (such as a standard door) and show how novel objects can be discovered and designed through processes of search, invention, and naïve physical reasoning. Key issues in this research are:

- The role of spacial knowledge and imagination.
- The representation of mechanical devices at varying levels of abstraction.
- Strategies of creativity, such as feature permutation and analogy.
- Heuristics for determining novelty of invention.
- Control of combinatorics.
- Organization and modification of episodic memory and how episodic memory affects processes of design and invention.

Participants: Jack Hodges (contact), Michael Dyer, Margot Flowers.

References: M. Dyer & M. Flowers (1984).

The GATE Project: Graphical AI Tools Environment

GATE is an integrated set of artificial intelligence development tools for the T language, a dialect of Scheme, running on Apollo Domain workstations. GATE consists of WEBS, a graphics knowledge representation system; DELON, a demon programming language; and TLOG, a logic (Prolog-style) programming language. All are available within a single interactive graphics interface. GATE also contains other AI tools including Flavors, D-nets, and a spreading activation mechanism.

Contacts: Erik Mueller (WEBS); Uri Zernik (DELON); Scott Turner (TLOG). *Other Participants:* Charlie Dolan, Michael Dyer, Margot Flowers, Seth Goldman, Alex Quilici, Walter Read.

C. Dolan & M. Dyer (1985b), E. Mueller & U. Zernik (1984), W. Read & M. Dyer (1985).

Current Research Projects at the UCLA Cognitive Systems Laboratory

Research at the Cognitive Systems Laboratory is focused on the theoretical foundations of reasoning and heuristics. Aspects of human reasoning which defy theoretical explanation motivate the formulation of new theoretical dilemmas and, conversely, computational principles of efficiency and resource management are hypothesized as explanations for prevailing patterns of cognitive behavior.

The research program comprises a number of research projects, directed by Professors Pearl, Parker, and Korf and concentrating in plausible and heuristic reasoning, relaxation methods in logic, and learning, respectively.

The Analysis of Heuristic-Search Algorithms

Our main concerns are to quantify the tradeoffs between the quality of the heuristics used, the complexity of the algorithms which they help guide, and the quality of the solution which they produce. Our results thus far have been in the areas of path optimization and game playing. Current efforts focus on greedy algorithms and constraint-directed backtracking which are more representative of human search strategies.

Participants: R. Dechter, G. Michon, I. Roizen, J. Pearl (contact).

References: J. Pearl (1984), R. Dechter & J. Pearl (1985b).

The Mechanical Discovery of Heuristics

We pursue the paradigm that heuristics are discovered by consulting simplified models of the problem domain. This

approach involves three major steps: (1) Simplification, (2) solution, and (3) advice generation. These are being formulated and implemented in the context of constraint-satisfaction problems where the targets of simplification are tree-structured constraint graphs which are known to be backtrack-free. The number of consistent solutions in the simplified problems is then extracted as a figure of merit for the options pending in the original problem.

Participants: R. Dechter (contact), J. Pearl.

References: J. Pearl (1985), R. Dechter & J. Pearl (1985a).

Computational Models of Evidential Reasoning

The aims of this project are to develop a computational representation of plausible reasoning that would facilitate an efficient integration of data from multiple sources and generate a coherent interpretation for that data. Our current efforts focus on Bayesian networks; namely acyclic graphs in which the nodes represent propositions (or variables), the arcs signify direct dependencies between the linked propositions, and the strengths of these dependencies are quantified by conditional probabilities. A network of this sort can be used to represent the deep causal knowledge of a domain expert and turns into a model of self-activated memory if the links are used not merely for encoding factual knowledge, but also for directing and activating the data flow in the process of reasoning about this knowledge. We find that, in many cases, Bayesian networks offer an ideal architecture for fusing and propagating beliefs, for controlling production rules in expert systems, and for hardware implementation of massively-parallel inference engines.

Participants: N. Dalkey, M. Ben-Bassat, I. Roizen, H. Geffner, G. Wang, J. Pinto, J. Pearl (contact).

References: J. Kim & J. Pearl (1983), J. Pearl (1985).

Discourse Generation

The ultimate goal of this project is to generate intelligent and congenial discourse about a subject for which rich machine knowledge is available. Our immediate task is the understanding of Meta-Subject-Utterances.

When reading or listening to an explanation of a technical subject, we notice the presence of conversational expressions like: “however,” “as I have stated before,” “next,” or “generally speaking,” etc. These expressions are not included in the discourse merely for decorative purposes, but carry important information that the listener uses to speed up the comprehension process. Accordingly, we have devised a functional taxonomy of these expressions based on their role in the listener’s processes, and cast their meaning in a representation that facilitates the generation of fluent and cogent discourse. As a testbed for our ideas, we have selected the domain of high-school algebra, and we are implementing a system called FIGMENT for generating commentaries on the solution of equations.

Participants: I. Zuckerman (contact), J. Pearl.

Reference: I. Zuckerman & J. Pearl (1984).

Computational Models of Learning

In this area we address the task of acquiring causal models from raw sensory data by invoking fictitious constructs called “hidden causes.” Causal models are treated as trees of random variables where the leaves are accessible to direct observation, while the intermediate nodes—the hidden causes—serve to mediate constraints between the visible variables. We have found that under certain conditions the structure of such trees can be inferred uniquely and effectively from measurements taken on pairs of leaves. Currently we study procedures of constructing causal trees as approximations to complex relational structures.

Participant: J. Pearl (contact).

Reference: J. Pearl (1985b).

Problem-Solving, Planning, and Machine Learning

Since search is a completely general method of solving any problem that can be expressed in a problem space, the real issue is how knowledge can be brought to bear to improve search efficiency. This amounts to exploring the different types of knowledge that exist, such as heuristic evaluation functions, subgoals, macro-operator sequences, and abstraction spaces, and examining exactly how they improve problem solving performance.

One theory currently being explored attempts to unify the use of heuristic evaluation functions in single-person problems with their use in two-person games. The theory explains why heuristic functions improve performance, and characterizes the effectiveness of a heuristic in terms of local information in a problem space. This makes it much easier to predict if a given heuristic will be effective for a problem, and makes it possible to automatically learn useful heuristics for problems.

Another theory being developed tries to unify problem solving and planning behavior. This is done by explaining planning as problem solving in the presence of knowledge sources, such as subgoals, macro-operators, and abstraction spaces. The theory supports quantitative analyses of the effect of these sources of knowledge on problem solving performance, and allows one to predict the efficiency of a planning program based on the knowledge it has available to it.

Research on machine learning is focused on general mechanisms that allow these different kinds of knowledge to be automatically acquired in a given domain. The goal is to build systems whose performance on problems in a given domain automatically improves, as more problems are solved. Experiments have shown that both macro-operator sequences and heuristic evaluation functions can both be learned by search in the given problem space. Efforts are underway to extend this type of learning to

subgoals and abstraction spaces as well. The research described here is supported in part by an NSF grant.

Participant: Richard E. Korf (contact).

References: R. Korf (1985a,b,c).

RAPPORT: Relaxation and Pattern-Oriented Rule Programming

Relaxation is a general computational paradigm for finding values meeting a set of constraints. With relaxation, goals are reached by successively finding unsatisfied constraints and satisfying them. Any "out of kilter" constraints are brought back to "in kilter" status. This approach has a growing list of applications, ranging from machine vision to decision making.

Relaxation appears to be useful in programming in that it permits problems to be specified declaratively (as sets of constraints) without regard to the specific solution method used. Pattern-oriented rule-based languages such as Prolog seem amenable to relaxation methods, and the focus of the RAPPORT project has been on clarifying the use of relaxation and its relationship with pattern-oriented programming. This research is supported in part by a UC Micro grant from the IBM L.A. Scientific Center.

Participants: D. Scott Parker (contact), Paul R. Eggert.

Principles of Relaxation in Logic Programming Systems

The literature is full of instances of relaxation. This research attempts to characterize what kinds of problems are reasonably solvable by relaxation, and looks at datatyping problems in logic programming languages. The problem of polymorphic type inference in such languages can be expressed with inequalities describing what is known about the argument types for each predicate. These inequalities are iteratively solvable using relaxation.

Participants: Ching-Tsun Chou (contact), D. Scott Parker.

Subproblem Relaxation

In subproblem relaxation, unsatisfied constraints are satisfied by (temporarily) ignoring them and considering the subproblem presented by the other constraints. Solutions obtained from the subproblem are used as initial approximations to solutions for the original. Applications of subproblem relaxation in mathematical programming and general pattern-matching processes are being investigated.

Participants: Kam-Pui Chow (contact), D. Scott Parker.

Inexact Reasoning by Constraint Propagation

Relaxation can also be applied in the modeling of inexact reasoning in expert systems. The area of reasoning under uncertainty has recently received a lot of attention. The emphasis here is on modeling the propagation of belief in an inference network as a constraint satisfaction process,

solvable by relaxation. The network is permitted to grow, and the propagation is not limited in direction. The issues of consistency checking and hypothetical reasoning are emphasized as central parts of a knowledge-based system.

Participants: Koenraad Lecot (contact), D. Scott Parker.

Current Research Projects in the Robotics and Vision Laboratory

Computer Vision Projects

Hierarchical image analysis involves the idea that algorithms can derive useful data by dealing with multiple resolutions of the same image. These methods use a set of increasingly-fine views to improve the speed of traditional image processing algorithms. Since that approach depends on a data structure method called quadtrees, some of our work has involved development of minimal (and hence, efficient from the view of program running time) image-storage versions; while other parts have focused on the utilization of the stored reduced-scene by algorithms that employ multiple levels of information. These methods are closely linked to design of procedures for acquisition of a part, as in factory automation, from a view of a scene.

Models from several images of a scene, indexing methods for storing and retrieving satellite imagery, and the design and manufacturing realm, where dynamic analysis (time evolution, motion of objects) is important. The concern of our work is with labels that are natural for humans in describing an element of a scene (as in a part seen in a machine drawing), and procedures that are useful for computers. Our work includes past accomplishments such as the creation of a working image database management system capable of simultaneously handling picture, map, and numeric data.

There is both a long history of image analysis from outline-drawings derived by computer processing of digitized scenes, and an equally deep tradition in sensor domains (IR-optical, radar, sonar, medical imaging: *e.g.*, ultrasound, computerized axial tomography, magnetic resonance imaging, etc.) of computing at a level that is much closer to the underlying objects. Essentially what we are doing in this work is in the latter tradition. In many situations what is discerned is not a viewable scene, but diffuse blobs; in others, a multitude of points and lines may be very clear but their very number obscures effective utilization. The project involves finding geometric objects by analysis of point data derived from their corners. We use minimum spanning tree and clustering methods from pattern analysis as elements of a new methodology that we are developing.

Participants: Allen Klinger (contact), Dariush Keyani-Yazdi.

Machine Vision: Sensing of Edge Motion. This project examines the use of a continuously moving image to provide dynamic edge determination. The first processing step beyond the initial image transduction is described in a forthcoming paper in the IEEE Transactions on Pattern Analysis and Machine Intelligence. The basis of a locally autonomous unit—referred to as a “triad”—has been developed that contains three equidistant intensity-step detectors. The triad can roughly determine edge direction and observed velocity from just the time ordering of the three cells created by a moving edge. More precise edge determination is possible via the timing relations among the three cells in a triad. This approach is different in many ways from more classical image analysis. Significantly, the method theoretically makes time continuous and distance discrete; the opposite is true of most vision systems that utilize discrete image frames. This eliminates the “correspondence” and other problems by implicitly tracking edges with the state of the triad unit.

Triad units are more sensitive to noise than are individual step detectors because several detectors comprise a single triad. This problem is quite significant for natural, noisy images. Current work is twofold. On the first part, a paper is in preparation which describes the noise sensitivity of triads. This work will provide a strong basis of evaluation for mechanisms that seek to reduce the noise sensitivity of the triad tessellation. The second part of the current work is developing a geometric unit that exploits local consistency conditions among adjacent triads to significantly reduce the noise sensitivity within the triad array. The goal of our current work is to produce a truly dynamic (*i.e.*, operation upon a continuously moving image) edge detector which provides the orientation and observed motion of a local edge, yet it must have a superior signal-to-noise ratio, extreme sensitivity, and it must be locally computable (*i.e.*, “inherently parallel”) and easily render itself to digital implementation. Simulations are being developed to allow experimental evaluation of the methods developed by this work.

The ultimate goal of this research effort is to develop a real-time machine that can formulate a concept of what objects are represented within an image without prior knowledge of the image domain.

Participants: Philip Kahn (contact), Jacques Vidal.

Expert Systems and Architectures for Realtime Processing

XRAM: An Expert System for Realtime Anomaly Management. A new representation of procedural logic, initially developed in J. Helly’s Ph.D thesis in the context of Space Shuttle malfunction procedures is the basis for the development of a realtime expert system supporting high performance aircraft flight operations. The research addresses the interface of telemetry data to the expert system, and localization of the search space within the knowledgebase

of the expert system as a function of flight operation phase and vehicle state. The system will serve as a testbed for the identification of performance criteria affecting the distribution of control and data between ground and airborne systems, between hardware and software and between conventional and special purpose parallel hardware, using programmable logic arrays (PLAs).

As air and spaceborne vehicles incorporate greater functional capability through the integration of increasing numbers of subsystems, the operation and maintenance of these vehicles grows exponentially more difficult. The addition of each new subsystem, itself usually with a sizable number of achievable states, multiplies the number of previously achievable vehicle states, and concomitantly the number of human experts necessary to monitor and operate the vehicle. The problem of accurately characterizing the vehicle state under nominal conditions is exacerbated in the presence of anomalies. Anomalies commonly introduce previously undiscovered states resulting from the interdependence of vehicle subsystems. In principle, such interactions can be anticipated although, in practice, the number of states limits the completeness of any systems analysis (Helly 1984).

As in our previous work on Space Shuttle subsystem anomalies the system will follow the PLA model of a knowledgebase. This representation is a data independent description of production rules, that can be directly translated into compact hardware. With this approach, the search procedures for anomaly detection and resolution and the subsequent processing could be performed entirely on-board or be distributed between airborne and ground-based processors. Current work is directed toward extending the approach of (Helly 1984) to incorporate temporal affects *i.e.*, to implement the notion of partial failures as time dependent functions of the vehicle state and a priori knowledge that can provide earlywarning of incipient failures.

Initial work is based on two advanced “fly-by-wire” aircrafts, the forward-swept wing X29A and the AFTI/F16.

Participants: John Helly (contact), Jacques Vidal

Reference: Helly (1984).

Learning in Boolean Networks. The goal of this research is to develop and study learning strategies for a class of networks consisting of simple programmable Boolean modules.

The networks considered consist of a number of identical functional nodes, each having a small number of inputs, and each implementing a Boolean function which is adjustable by a local miniature finite state automaton. The input and output leads of the nodes are interconnected in a hierarchical fashion and the resulting network implements a larger (adjustable) Boolean function. This class of networks includes (but is certainly not limited to) Perceptron-like systems and many other threshold gate networks, but

we are mostly concentrating on more general networks of digital nodes (which in theory could be made functionally complete).

The object of the research is to find strategies for "teaching" this Boolean network a desired functional behavior without the user having to deal with the specific internals of the network. In other words, the ideal learning system should consist of a set of small processes running concurrently, each controlling one node, and each utilizing as little local memory or internode communication as possible.

Participants: Rik Verstraete (contact), Jacques Vidal.

Current Research in Robotics and Intelligent Manufacturing

Development of an action/decision system for operating an automatic factory This project is aimed at permitting factory managers to define relatively easily a computer program capable of performing the necessary decisions and initiate the actions required to run an automatic factory. This system is being developed with a graphical user interface, simulation capabilities and the use of PROLOG (a language used in Artificial Intelligence) as the language in which the necessary rules are going to be expressed.

Smart CAD Systems. This project consists in providing mechanical CAD systems with the capability to monitor a user designing a part or a system, checking his work for correctness, manufacturability, economy and other desirable features as required. Such a system will also eventually allow drafting departments to dispense with human checkers.

Development of hardware and software systems aimed at producing intelligent robots provided with sensory and feed-back capabilities. This system should allow the user to specify desired robotics task in a very high level language and permit him to predict the resulting actions of the robot through graphical animation.

Automation of various manufacturing planning tasks including: NC (numerical control) programming, NC program verification, and manufacturing process planning by means of expert systems.

The CAD lab operates some 30 graphics consoles controlled by an IBM 4341 computer donated to UCLA last year by IBM under IBM's CAD/CAM university grant program. The CAD projects are sponsored by IBM, Lockheed, Hughes, and Northrup.

The robotics lab currently houses three commercial robots and a vision system including an IBM 7565 and IBM 7535. The robotics project is sponsored by a grant from IBM.

Contact: Michel Melkanoff.

References

Alvarado, S J, Dyer, M. G & Flowers, M (1985) *Memory representation and retrieval for editorial comprehension*. Proceedings of the

- Cognitive Science Society Irvine, Ca
 August, S E & Dyer, M. G. (1985) *Understanding analogies in editorials* IJCAI-85.
 Dechter, R , & Pearl, J. (1985a) *The anatomy of easy problems: A constraint-satisfaction formulation* IJCAI-85
 Dechter, R , & Pearl, J (1985b) Generalized best-first search strategies and the optimality of A*. *JACM*, July, 1985.
 Dolan, C P (1985) Memory based processing for cross contextual reminding UCLA Computer Science Dept M.S. Thesis
 Dolan, C P. & Dyer, M. G. (1985a) *Learning planning heuristics through observation*. IJCAI-85
 Dolan, C., & Dyer, M G (1985b) The world's shortest Prolog compiler? Tech. Rep UCLA-AI-85-2
 Dyer, M. G. (1983a) *In-depth understanding* Cambridge, MA: MIT Press
 Dyer, M G (1983b) The role of affect in narratives *Cognitive Science* Vol 7.
 Dyer, M G (1981) The role of TAUs in narratives *Proceedings of Cognitive Science Society* Berkeley, CA
 Dyer, M. G. (1983c) *Understanding stories through morals and reminders* IJCAI-83.
 Dyer, M. G. & Flowers, M (1984) *Automating design invention* Proceedings of AUTOFACT-6, sponsored by Computer and Automated Systems Association of Society of Manufacturing Engineers.
 Dyer, M. G., & Flowers, M (1985) *Toward automating legal expertise* Proceedings of the First Law and Technology Conference. Houston, Texas
 Dyer, M G & Lehnert, W G. (1982) Question answering for narrative comprehension In LeNy & Kintsch (Eds), *Language and Comprehension* North-Holland
 Dyer, M G. & U. Zernik. (1985) *Parsing strategies and language learning* Proceedings of AI-1985 Conference Long Beach, CA.
 Feifer, R G , Dyer, M G , & Flowers, M Teaching inferencing strategies. Tech. Rep UCLA-AI-85-16
 Flowers, M (1982) *On being contradictory* AAAI-82.
 Flowers, M & Dyer, M. G. (1984) *Really arguing with your computer in natural language*. NCC-84
 Flowers, M , McGuire, R & Birnbaum, L (1982) Adversary Arguments and the Logic of Personal Attacks In Lehnert & Ringle (Eds.), *Strategies for Natural Language Understanding* Trenton, NJ: LEA Press
 Gasser, M (1985) Generating sentences in a second language. UCLA AI Lab Tech Rep
 Goldman, S R , Dyer, M G & Flowers, M. *Precedence-based legal reasoning in contract law* IJCAI-85
 Helly, J J , Jr (1984) A distributed expert system for space shuttle flight Control PhD thesis, UCLA
 Kim, J , & Pearl, J (1983) *A computational model for combined causal and diagnostic reasoning in inference systems* IJCAI-83, 190-193
 Korf, R E. (1985a) Depth-first iterative-deepening: An optimal admissible tree search To appear in *Artificial Intelligence*
 Korf, R E (1985b) *Learning to solve problems by searching for macro-operators* London: Pittman Publishing
 Korf, R E (1985c) Macro-operators: a weak method of learning To appear in *Artificial Intelligence*
 Mueller, E T. & Dyer, M G (1985) *Daydreaming in humans and computers* IJCAI-85
 Mueller, E T & Dyer, M G (1985) Towards a computational theory of daydreaming Proceedings of the Cognitive Science Society Irvine, CA
 Mueller, E T., & Zernik, U (1984) GATE reference manual. Tech Rep UCLA-AI-84-5
 Pazzani, M J (1985) Explanation and generalization based memory Tech Rep UCLA-AI-85-13
 Pearl, J (1985a) Bayesian networks: A model of self-activated memory for evidential reasoning Tech Rep R-43, Cognitive Systems Laboratory