AI Magazine Volume 6 Number 3 (1985) (© AAAI)

RESEARCH IN PROGRESS

Artificial Intelligence at MITRE

Joseph L. Katz

The MITRE Corporation, Post Office Box 208, Bedford, Massachusetts 01730

The MITRE Corporation is a scientific and technical organization engaged in system engineering activities, principally in support of the United States Air Force and other government agencies, and primarily in the field of information systems. MITRE is a special kind of engineering organization. The corporation is a Federal Contract Research Center, a designation covering the handful of independent institutions that perform government-sponsored research. It is an independent, nonprofit corporation designed and managed to provide long-term assistance to government agencies in planning, design, procurement, and testing of their information systems.

Artificial Intelligence (AI) research at MITRE started in the early 1970s with the development of MATHLAB by Carl Engelman. MATHLAB was an on-line system providing a computer aid for the mechanical symbolic processes encountered in mathematical analysis. Many of the ideas generated in its development have influenced the architecture of later systems. Systems such as IAM, MACSYMA, and REDUCE have acknowledged conceptual debts to MATHLAB; ALADIN, SCRATCHPAD, and SIN (the Moses Symbolic INtegration Program) incorporate the actual code of some of the more important MATHLAB subsystems.

MITRE has continued artificial intelligence research through internally funded independent research and development (IR&D) programs. In 1978, the Directorate of Mathematical and Information Sciences began sponsorship of a MITRE program to investigate and demonstrate the applicability of promising artificial intelligence technology. The project came to be known informally as "KNOBS," an acronym for Knowledge-Based System. Subsequently, Rome Air Development Center took over support of the project and continues to fund part of our AI research effort.

MITRE's current research is summarized below. The Bedford center is supported by 15 Symbolics Lisp machines netted to two Vax-780 file servers, while the Washington center is supported by both a classified and an unclassified facility, with 2 Lambdas and 2 Symbolics Lisp machines respectively netted to Vax-780 file servers. Both centers support creative groups of people who generate exciting new ideas.

MITRE Bedford AI Programs

Planning and Reasoning

Research into planning and reasoning started with the development of the KNOBS system. KNOBS was an early effort to unite several capabilities: Frame representations for data, rule-based inference, English understanding using a conceptual-dependency parser for question-answering and control, and a constraint-based approach to planning. This complex architecture evolved from a simple rule-based system as a result of the functional and performance demands on the system.

Our current research is directed towards expanding the capabilities in KNOBS to include replanning through an implementation called the KNOBS Replanning System (KRS). The domain dictates that the system create plans for complex resource reallocations, revise those plans when expectation failures occur, and also minimize the impact of any necessary changes on the remainder of the plan. In replanning, the system must deal with previously constructed plans as objects it can reason about and alter; our research has focused on mechanisms to allow for meta-level

For more detailed information about any of MITRE's projects please contact Joseph Katz (KATZ@MITRE-Bedford) or Richard Brown (LINUS!BROWN@MITRE-Bedford) at the Bedford center or Peter Bonasso (BONASSO@MITRE) at the Washington center reasoning and planning. In addition, we have recognized the importance of global strategies to the planning process and are exploring the elements of a system that can use global strategies in its planning processes, understand its own use of strategies, and explain these uses.

To accomplish these goals, we have focused on the development of a declarative formalism for expressing meta-planning knowledge. The formalism under development will be able to express global strategies and heuristics along with control information such as subprocess orderings, possible concurrency, and interprocess communication. It will also be able to accommodate widely diverse domains and be interpretable both procedurally and declaratively so that the system can intelligently explain its own activity.

We have partially developed and implemented a declarative formalism for describing planning strategies in KRS to reason about resource reallocation. Continued development will focus on extending the formalism to deal with problems of dynamic replanning (replanning during plan execution) and to replace certain procedural problem solution techniques within KRS.

Principal Investigator: Stuart Goldkind, Ph.D.

A Relational Grammar-Based Parser for Interacting with Expert Systems

The King Kong Parser development is aimed at building a portable natural language interface for expert systems of the kind typified by the interface in KRS. The King Kong Parser depends upon the theory of Relational Grammar (RG) to provide a bridge between its syntactic and semantic components. The former is a syntax analyzer similar to that developed by Marcus, and the latter is a semantic net associated with a set of scripts and inference rules dependent upon the domain of application. Relational grammar deals with language universals such as subject, object, and instrument that are, in a sense, midway between syntax and semantics. We hope to make our parser truly portable, so that it could even be retargetted to another language with a small amount of effort. Since the parser depends on language-independent categories supplemented by a self-contained set of language-dependent syntactic rules, such portability is a real possibility.

The semantic goals are to enable the parser to handle simple metaphorical word extension, such as the use of synecdoche and certain types of metonymy. We also hope to build in some learning capability to allow for "chunking" of concepts and derivation of words across parts of speech.

A final goal of our parsing system is to integrate it into an intelligent interface package that would be able to discern user goals and to regulate information flow across linguistic and graphic media accordingly.

Principal Investigator: Candace E. Kalish.

Knowledge-Based Automatic Programming

Automatic programming is a term devised in the 1950's to describe the process of using an algebraic compiler. Such compilers liberated application programmers from concern with the binary nature of machines and with the tedious allocation of registers. Exploitation of knowledge-based technology to allow the inclusion of both application-specific facts and programmer expertise yields new capabilities for the automation of programming. Previous work in knowledge-based automatic programming has developed much of the theoretical basis needed to mechanize many aspects of the programmer's art: representation selection, data-structure design, reasoning about and managing side-effects of database updating, discovering clever adaptations of algorithms, and using programming "tricks." A partial implementation of this theory has been demonstrated by this project in the context of a simple color graphic situation map database and display system. A near-term goal is to produce a more mature automatic programming system that can, in several different applications, be used by the relatively inexperienced to create, test, and modify software systems. Achieving this goal will involve: addressing "ease of use" issues, further developing and evaluating the concept of automating the utilization of problem-solving strategies for controlling combinatorial explosion, and extension of automatic programming technology to include monitoring the execution of code produced.

Principal Investigator: Richard H. Brown, Ph.D.

Knowledge-Based Query Optimization

The object of this project is to investigate the applicability of knowledge-based automatic programming technology to generating fast relational database queries using domain knowledge. Modern database systems have a powerful, high-level query language for extracting information from a database. This leads to systems that are easy to use, but they can be slow The goal of query optimization is to find alternative methods for answering the queries to decrease the time required to retrieve the information. Existing optimizers never make use of the meaning of the information stored in the database; instead they look at the interrelations of the data. Human experts can often improve on the results of optimizers by using knowledge of the meaning of the data to evaluate the cost of a method for answering a query, as well as to generate new alternatives. This fact suggests that the application of artificial intelligence techniques is appropriate. Since query optimization is closely related to writing programs, we plan to apply the mechanisms developed by the Knowledge-Based Automatic Programming project to query optimization.

Principal Investigator: John D. Ramsdell, Ph.D.

Scheduling Against Nonlinear Constraints

Knowledge-based and expert planning systems employ many constraints to generate, check, and criticize proposed plans. The present expert systems approach to scheduling—and more generally the "optimal" allocation of all relatively scarce resources—is typically an *ad hoc* variant of "generate and test," which searches through all possible solutions until one is found that violates no constraints. This process is intolerably slow and usually does not find the minimum cost solution first. Using standard techniques from the field of operations research (*e.g.,* "linear" and "dynamic" programming) is problematic because the actual constraints of the planning problem are often nonlinear and even nonnumerical. Our approach involves a technique based on "reasoning by analogy." Given a planning problem involving resource allocation, our approach uses heuristics to generate several linear subproblem analogs, each of which can be solved using techniques like linear and dynamic programming. Other heuristics will be devised to then incorporate those analog solutions back into the original, nonlinearly constrained planning problem.

Principal Investigator: Richard H. Brown, Ph.D.

Future Generation Computer Architectures

Revolutionary improvements in computation system capabilities are suggested by technological advances in artificial intelligence and related fields. However, the full realization of these improvements will ultimately require an increase of several orders of magnitude in computer processing speed. To realize these speed increases, architectures comprising thousands of processors with shared access to an equally large number of memory units will be required. Unfortunately, conventional methods of achieving the required interconnections prove intractable for more than a few processors. In fiscal year 1984, we designed a computer architecture where a "smart memory" (in which each cell has an associated computation capability) provides both the required interconnection and a significant portion of the system's computing power. We also constructed simulations and assembled a test bed in which portions of the architecture could be emulated. We are refining the design to investigate software implications of the architecture, to evaluate its suitability for artificial intelligence applications, and to construct a portion of the architecture in very large-scale integrated electronics.

Principal Investigator: Edward H. Bensley.

EMPRESS—
Expert Mission Planning and Replanning Scheduling System

The objective of EMPRESS is to develop a prototype expert system for the National Aeronautics and Space Administration (NASA) to use in scheduling activities related to preparing cargo payloads for space shuttle missions. EMPRESS will provide two forms of support: first, the system checks schedules by tracking resource usage and identifying conflicts, and by comparing allotted task duration to *a priori* temporal expectations or user stated scheduling heuristics; second, the system attempts to propose partial solutions to scheduling problems. This second capability—automatic replanning—distinguishes the AI approach from conventional approaches. These proposed solutions will attempt to retain the rationale and essential user- and system-provided constraints on tasks.

At no time will the system relieve the user of the responsibility of independently insuring that the schedule is correct; rather, it will detect difficulties with some degree of reliability and propose potential solutions to scheduling problems.

Principal Investigator: Joseph L. Katz, Ph.D.

KNEECAP—Crew Activity Planner

The KNEECAP system is currently under development to demonstrate the applicability of using expert system technology for NASA crew activity planning aboard the Space Shuttle. The system is derived from the architecture used in KNOBS, which was constructed to aid air mission planning. KNEECAP is similar to KNOBS and has facilities to help the user plan mission items interactively by drawing on a well-structured knowledge base. Constraint checking is performed after the user proposes a value for an element of the plan and failures are indicated to the user with some explanation. Acceptable choices can be enumerated and ordered for a plan element based on existing choices of other plan elements. The user can invoke the autoplan facility which finds a consistent set of preferred values for all of the remaining plan elements if possible. Limited automatic replanning is implemented by using a forward-chaining rule interpreter to initiate replanning when changes to the database invalidate the current plan. Such changes might occur when a piece of equipment is damaged, a crew member becomes ill, or unexpected delays occur in completing the previous task.

Principal Investigator: Judah Mogilensky.

LES—LOX Expert System

Loading the Space Shuttle's Liquid Oxygen (LOX) propellant into its external fuel tank is performed through a sequence of some 10 major steps in a fixed script, controlled by the highly automated Launch Processing System (LPS). In the first application of artificial intelligence sponsored by the Kennedy Space Center, an experimental knowledge-based system called LES has been constructed to aid in monitoring the LOX loading process and in diagnosing faults. LES accomplishes this by reasoning directly from knowledge of the LPS' structure and the control relationships between its components. LES' knowledge base includes most of the LOX portion of the Launch Processing System at KSC, including analog and discrete commands and sensors, and other objects (transducers, relays, solenoids, valves, LOX pressures and temperatures, etc.) whose state the system is designed to control or monitor.

Intuitively appealing diagnostic and knowledge representation techniques were developed through consultation with system engineers. Control relationships that determine the consistency of sensor measurements are represented by symbolic expressions embedded in frames and are used for both fault detection and location. These functional relationships are stored in exactly one place, so they must be inverted to determine hypothetical values for possibly faulty objects; the process has been extended to include conditional relationships. A demonstration system is now operating, and later evaluation will use live launch data.

Principal Investigator: Ethan A. Scarl, Ph.D.

Heuristics Research

While current progress in artificial intelligence is generally encouraging, the once promising field of automatic theorem proving has been relatively less successful. A particularly neglected basic research topic has been the guessing procedures, called heuristics, on which nearly all theorem proving programs rely. Because exhaustive search for proofs, no matter how cleverly implemented, is not feasible, most provers call upon heuristics, usually to pick a "most likely true" alternative from a given selection of assertions. Fundamental to this project's approach is a new extension of an old idea from symbolic logic. It explains the meaning of a formal statement in terms of a contest between two sides, similar to a debate. This allows certain techniques from the somewhat isolated field called adversary analysis to be applied for the first time to automatic deduction. We will separate out heuristics for direct empirical and theoretical analysis and will attempt to incorporate several useful aspects of the way humans actually reason about mathematical problems into high quality heuristics oriented towards specific types of problems. Results will benefit automatic theorem provers, software verification efforts which depend on formal proofs, and possibly the average-case performance of a variety of mathematical algorithms.

Principal Investigator: Leonard G. Monk, Ph.D.

MITRE Washington AI Programs

The ANALYST Project

The ANALYST project grew out of a sponsored effort to apply production system techniques to multi-sensor analysis in a military context. In recent years the system has grown into a general knowledge-based aid for intelligence staff personnel. During its evolution ANALYST has served as a testbed for production system research issues involving spatial and temporal techniques and evidential reasoning. More recently, due to an operational experiment under DARPA auspices, the research staff has been motivated to delve into knowledge-acquisition issues and machine learning using ANALYST as the backdrop system (see below).

ANALYST uses multiple reasoning paradigms which interface a data-driven matcher of syntactic patterns of sensor activity with a goal-driven propositional system for establishing trends and making predictions. Data and knowledge structures are all frame-based, and spatial structures are organized into quadtrees for efficient search of map data. Research and "production" versions of the system are implemented in both Franz and Zeta Lisp dialects.

Principal Investigator: Pete Bonasso.

The Knowledge Acquisition Project

This project uses a three-pronged attack on the problem of getting rule based systems to learn from experience, that experience being first input from knowledge-engineering, then input from the environment, and then from introspection. The first prong attempts to make the system resistant to new rules by grounding a subset of its established rules in a particular proven behavior, *i.e.*, by endorsing outcomes known to be valid. The endorsements take the form of a subset of the world data at snapshots in time attached to the rule-set. New rules violating the endorsements will not be easily accepted into the system.

The second prong is to implement a version of Holland's bucket brigade technique to condition the rules over time. Weak or contradictory rules will tend to atrophy in their ability to contribute to the next inference period, and strong rules will grow in their capacity to contribute.

Finally, for automatic learning, we are investigating the application of various rule-modification schemes including those of Michalski, Lenat, and Holland, by recasting these schemes in the production formalism. The idea is, of course, to have the system hypothesize new rules to itself, especially when its behavior is consistently erroneous in some area.

The ANALYST system with its integrated forward and backward inferencing paradigms, and its spatial structures for capturing world scenes for endorsements will serve as the implementation vehicle.

Principal Investigator: Jim Antonisse

The Adaptive Simulation Project

We have used RAND's Rule Oriented Simulation System for several years to provide object-oriented models of the military environment as data feeds to the other research projects. The behavior-based simulation has proven effective in providing causal and relational data for production rules, semantic nets, and for spatial reasoning about terrain and world object collections. The current model is termed the Battlefield Environment Model (BEM).

There are distinct advantages to using the object-oriented approach to simulation modeling, particularly as concerns adapting the model, often on the fly, to various data requirements and more importantly to various behavior sets in the environment. But there are speed deficiencies as well. We are working with RAND, DARPA, and several industrial agencies on studying advanced machine architectures. Lisp machine versions are being examined as well as Intel's Hypercube and BBN's butterfly architecture.

Principal Investigator: Dick Nugent.

The Oplanner Project

OPLANNER was developed during 1984 to support a demonstration of a joint logistics planning aid. OPLANNER is an opportunistic planning system based in part on an earlier MITRE Washington system called KBS. In support of the logistics problem, OPLANNER develops partial plans with estimated resource requirements. OPLANNER is a hierarchical, rule based system that associates rule sets with nodes in the hierarchical plan structure. That is, as the plan is refined and expanded, rules and constraints are associated with each node in the plan network. The activity of the rule sets associated with each node may be individually controlled by other rule sets. This may be viewed as an active plan development network whose expansion may be differentially controlled. The nodes of the plan are instances of the basic planning primitives relevant to the domain of interest. Each node contains instance variables associated with that domain primitive. Instantiation of the nodes is done by rules responding to the goals, constraints and problems of the current planning problem. Domain-specific information is contained in a large semantic network.

OPLANNER was implemented on a VAX 780 in Franz Lisp and is currently being converted to the Symbolics. The underlying production system (SAPS—Still Another Production System) allows full use of Lisp in both antecedents and consequents. Efficiency measures detect changed values; however, a full discrimination network has not been implemented.

In addition to the ALLIES project, current effort is concentrated on a highly interactive, mixed initiative user interface and the efficient inclusion of multiple plan alternatives.

Principal Investigator: John Benoit.

The Allies Project

As a Federal Contract Research Center, MITRE has identified several areas of command and staff activity which may be helped by AI techniques. One of these is situation monitoring. The ALLIES (Air/Land Loosely Integrated Expert Systems) project is an attempt to set up a co-operative of knowledge-based systems which allows us to investigate situation monitoring techniques Basically we are having ANALYST and the BEM object-oriented simulation provide a data feed of the situation to the OPLANNER system. Since OPLANNER generated the original directives to actors in the BEM, this dynamic situation feed should trigger dependency mechanisms in OPLANNER to notice when preconditions and/or expected effects of the plan operators are being violated.

This project also surfaces research issues concerned with integrating knowledge-based systems. Since both O-PLANNER and ANALYST are doing some overlapping analysis, a division of labor must be determined. Questions such as whether there should be a common terrain representation and which system lexicon is semantically appropriate (an OPLANNER "task" is a BEM "behavior" and is seen by ANALYST as a "manifestation") will be addressed.

Since this project operates in a conflict environment (the European battlefield), we also plan research in counterplanning techniques and in using a combination of scripts and reverse plans to determine and predict opponents' intentions.

Principal Investigator: Pete Bonasso.