

Basic Artificial Intelligence Research at the Georgia Institute of Technology¹

*Ashok K. Goel, editor, with Ronald C. Arkin,
Kurt P. Eiselt, Janet L. Kolodner, Daryl T. Lawton,
and Ashwin Ram*

■ *AI research is conducted at a number of academic and research units at the Georgia Institute of Technology. Some of this research is basic in nature, and some has an applied character to it. This article briefly describes basic AI research in the College of Computing at Georgia Tech.*

The AI research group in the College of Computing currently consists of six full-time academic faculty members: Ronald C. Arkin, Kurt P. Eiselt, Ashok K. Goel, Janet L. Kolodner, Daryl T. Lawton, and Ashwin Ram. It also includes Richard Billington, a full-time research scientist, and about 20 graduate students. As a group, we cover a fairly large portion of the AI research spectrum, including the use of knowledge in problem solving, language processing, explanation, perception, and robotics; representation and organization of knowledge; and knowledge acquisition, learning, and instruction. Much of our research has a strong cognitive flavor. Some of it is inspired by cognition, and some explicitly models cognitive processes.

Problem Solving

Ashok Goel and Janet Kolodner conduct research on problem solving. Kolodner has long been investigating the use of case-based reasoning for solving a range of complex problems in a variety of domains. Goel is exploring the integration of different types of knowledge and methods of reasoning for planning and design problem solving.

Case-Based Reasoning

In *case-based reasoning* (Kolodner 1990), a reasoner solves new problems by remembering previous situations similar to the new situation.

Case-based reasoning can mean adapting old solutions to solve new problems, using old cases to critique new solutions, using old cases to interpret new situations, or using old cases to explain new situations. Kolodner's previous projects addressed issues in case-based reasoning in a variety of domains: mediation of resource disputes (MEDIATOR [Kolodner and Simpson 1989]), mediation of labor contracts (Persuader [Sycara 1987]), medical diagnosis (SHRINK [Kolodner and Kolodner 1987]), and meal planning (JULIANA [Shinn 1989]). In these projects, she and her students investigated issues such as the role of case-based reasoning, the control of case-based reasoning processes, the use of multiple cases, adaptation strategies, and the reuse of plans in a changing environment. These investigations led to a number of conclusions:

*This article briefly
describes basic AI
research in the College
of Computing at
Georgia Tech.*

First, in general, several cases are needed to solve complex problems. Second, the reasoner's goals have several important functions in case-based reasoning: They help to choose the means of indexing cases, select the best of several available cases, control the use of multiple cases, and focus on an appropriate part of a chosen case. Third, in case-based reasoning using multiple cases, there is a need to check the consistency of the proposed solutions. Although a

case-based reasoner can suggest solutions, other processes are needed to maintain consistency between the solutions and provide processing control. Fourth, because the world is not completely predictable, a case-based reasoner needs a means of interrupting its plans to respond to changes in the world, which requires a sophisticated goal scheduler. Fifth, the processes used in case-based reasoning are good for tasks other than case-based reasoning. The same adaptation heuristics that are used to adapt an old solution to a new situation are useful at execution time for adapting a plan to an unpredicted situation.

Kolodner continues to explore case-based reasoning in many of her ongoing projects. These projects address a number of issues. First, how can experience be represented and organized? Second, at what points in reasoning is experience used? How is it used at these points? What roles does it play in reasoning? Third, what kinds of performance differences can we expect in a novice reasoner, that is, one with little experience, as opposed to a relatively more expert reasoner that has considerable experience with novel cases? Fourth, what kinds of memory differences can we expect in comparing two such reasoners? Fifth, how does the evolution from novice to expert happen? Sixth, what retrieval strategies are necessary to access appropriate previous experiences? Seventh, based on what we know about the role of experience in reasoning, what kind of systems can we create that can better aid people in solving problems?

Case-Based Planning In the MEDIC project, Kolodner and Roy Turner (who wrote his Ph.D. thesis on this project) explored the use of case-based reasoning in planning (R. Turner 1989). Specifically, they investigated the reuse of old plans in a dynamically changing world where planning and execution must necessarily be interleaved. MEDIC's task domain is the diagnosis of pulmonary disorders. It reasons, as a doctor would, about how to carry on a diagnosis, for example, which symptoms or hypotheses to follow up next. Thus, diagnosis is addressed as a planning problem. Although MEDIC uses old diagnostic plans to guide its reasoning (its case-based component), much in the environ-

ment prevents it from carrying out its plans as previously done. Symptoms are not always discovered in the same order, many problems can simultaneously be investigated, test results obtained in a previous case might not match those obtained in the new one, and so on. Although parts of old plans are appropriate in solving the new problem, no old plan can be followed start to finish to solve the new problem. In general, the order of steps needs to be changed, and some steps need to be substituted or deleted. In the general case-based reasoning framework, there is an assumption that much of this adaptation can be done before execution. In MEDIC's domain, much of the adaptation must be deferred until some execution has been done. To enable this execution, MEDIC adds a sophisticated goal scheduler to the case-based reasoning framework. Solutions from old cases are added to the goal scheduler subgoal by subgoal. Subgoals are prioritized for the new situation. Each time a new subgoal is chosen (at execution time), the highest-priority one is chosen and adapted to the new situation. To make this approach work, MEDIC has to specify the means of determining goal priorities and the means of representing cases so that their parts can appropriately be accessed. Two ongoing projects that build on MEDIC—CELIA and Expeditor—are described in the section on learning.

Case-Based Design In the JULIA project, Kolodner and Thomas Hinrichs (who is writing his Ph.D. thesis on this project), explore the role of case-based reasoning in design (Hinrichs 1988, 1989, 1990). JULIA's design domain is meal planning. Although a commonsense type of task, meal planning shares much with other design domains. Problems are specified in terms of functional constraints, and solutions are descriptions of concrete artifacts. Problems are underconstrained and often ill defined. They are too large to solve as one chunk and need to be solved in parts, but the parts interact with each other in strong ways.

Beginning with a functional specification of a problem, JULIA uses previous cases to suggest solutions and warn of potential problems. Its adaptation strategies can adapt an old solution to fit a new situation (for example, it can turn lasagne into vegetarian lasagne if vegetarians are

Goel's research on problem solving focuses on knowledge-based planning and design problem solving.

invited for dinner). Kolodner and Hinrichs's emphasis in JULIA is on two problems: a functional architecture that allows a case-based reasoner to function in a complex, real-world domain and a concrete theory of adaptation. The architecture she is working on incorporates components for the breaking of a problem into parts, the maintaining of consistency between the parts, knowledge and case access, and adaptation. Concretely, a constraint propagator maintains relationships between parts of a problem, a reason maintenance system notices inconsistencies and points them out, a problem decomposer breaks problems into parts, a case memory organizes and retrieves cases and general knowledge, and an adaptation engine adapts old solutions to fit new situations and fixes (patches) proposed solutions to better solve problems. Together, the case memory and adaptation engine make up the case-based reasoner. The combination of constraint propagator and case-based reasoner allows highly underconstrained problems to be solved. The case-based reasoner proposes and adapts solutions and warns of potential problems (allowing early commitment to solutions when experience can give guidelines), and the constraint propagator, combined with the reason maintenance system, maintains relationships between parts of the problem (allowing late commitment when it is necessary) and notices places where the proposed solution is not right and requires adaptation.

Integration of Multiple Methods of Reasoning

Goel's research on knowledge-based planning and design problem solving. His work emphasizes the integration of different types of knowledge and methods of reasoning.

Knowledge-Based Planning In his earlier work, Goel studied political decision making by nation-states viewed as cognitive agents. This work led to the development of a computational model in which the decision-making behavior of nation-states is viewed as a kind of reactive planning (Sylvan, Goel, and Chandrasekaran 1990). The conditions under which this model is applicable pertain to the internal structure of the nation-state and the nature of the external events. The JESSE system simulates the reactive planning model for Japanese decision making in the domain of energy supply security. In JESSE, the task of reactive planning is decomposed into several subtasks: matching event data with stored concepts, refining the activated concepts, selecting stored plans, and refining the selected plans. JESSE uses several different types of knowledge, for example, knowledge of domain concepts, constraints, and plans as well as knowledge for matching, selecting, and refining concepts and plans. The stored concepts and the stored plans are organized in generalization hierarchies. The concepts and the constraints act as functional indexes to the plans. JESSE has been extensively and successfully evaluated against the historical record of Japanese actions concerning energy supply security.

Knowledge-Based Design In more recent work, Goel studied the design of physical devices such as electric circuits and heat exchangers. This work led to the development of a computational model of innovative design that integrates case-based and model-based reasoning (Goel 1989, 1990; Goel and Chandrasekaran 1989b, 1990). In the case-based approach to design, a novel problem is solved by adapting a design known to solve a related problem. Adapting a known design to solve a related problem by the commonly used methods of heuristic association and search can be computationally costly, however, if the adaptation search space is not small. Thus, the adaptation space needs to be decomposed into smaller and simpler spaces that can be searched more efficiently and effectively. The knowledge for decomposing the adaptation search space can be represented as a behavior structure model that specifies how the structure of the known design results in its output behaviors.

The model aids case-based design in several ways:

First, it identifies conceptual primitives for indexing the known designs stored in a case memory. Second, it identifies elementary types of behavior transformations and elementary types of structure modifications (Goel 1990). Third, it provides knowledge for decomposing the adaptation search space into smaller spaces so that the search for the needed structure modifications is localized (Goel and Chandrasekaran 1989a). Fourth, it leads to a novel method for simulating the behavioral effects of structure modifications. In this method, the output and causal behaviors of the modified design are derived by revising the output and causal behaviors of the known design (Goel 1989).

This computational model for adaptive design problem solving was simulated in the KRITIK system. KRITIK uses several different types of knowledge, including knowledge of design cases and device models. It unifies the methods of case-based reasoning, model-based reasoning, heuristic association, heuristic search, task decomposition, object decomposition, plan instantiation, and constraint propagation into a single architecture for design problem solving. KRITIK was successfully evaluated in the domains of simple electric circuits and heat exchangers.

New Projects

Three new research projects on problem solving are being pursued. In one project, Kolodner is investigating creative problem solving (Kolodner and Penberthy 1990). In this project, she explores augmentations to case-based reasoning that result in innovative or creative problem solving. Based on observations of people engaged in creative problem-solving activities, she discovered that one of the major activities creative problem solvers engage in is the exploration and evaluation of alternatives, often adapting and merging several possibilities to create a solution to the new problem. The reasoner starts with a partial, incomplete problem specification and, through a series of example retrievals and evaluations, eventually defines the problem more clearly and creates a solution. Several solutions are considered, and the final one has elements of many that were considered before it.

Kolodner addresses several issues in

this context. First, she is exploring a means of making retrieval processes more flexible because brainstorming requires being open to remembering things that might not be relevant on the surface. Second, she is investigating evaluative processes. Evaluation is a key to exploration because its results are used to both further define a problem and create a solution. Of particular importance is focusing evaluative processes so that appropriate evaluations are done. Third, she is exploring a means of merging parts of several solutions to create a new one. Fourth, she is exploring the ways in which problem specifications

The process of language understanding can be viewed as a series of decisions

can be changed over time as a result of evaluating opportunities for solutions. Many times creative solutions arise as a result of redefining problems. The program we've implemented explores its library of alternatives to propose ways of using up leftover white rice.

In another project, Kolodner and Goel are investigating the role of case-based reasoning in architectural design. In this project, we are coupling our previous work on design problem solving with guidelines for building architectural case libraries to create a system that could function as an architect's design assistant. The proposed system begins with a library of architectural cases in which the cases are indexed on features provided by expert architects. An architect using the system describes his(her) problem to the system, and the system recalls similar cases. Recalled cases will play several roles. They can warn of the potential difficulties in solving the problem; they can suggest methods for solving the problem or some part of the problem; they can suggest methods for adapting or repairing a proposed solution; and they can project the outcome of a proposed solution, thus helping with solution evaluation. This work is being done in collaboration with researchers at the School of Architecture.

In a third project, Goel is investigating the integration of model-based and case-based methods for planning routes in physical spaces. This project led to the development of the ROUTER family of path-planning systems. ROUTER1 uses a topographic model of the Georgia Tech campus to plan routes from one intersection on the campus to another (Goel et al. 1991). ROUTER2 plans new routes by retrieving and adapting previous route-planning cases. ROUTER3 integrates the model-based and case-based methods for route planning. The control architecture for ROUTER3 is especially noteworthy. It uses a processing control in which the model-based and case-based methods compete with each other to solve a given task or subtask. The selection of a specific method depends on the state of problem solving and the degree of match between the knowledge needed to solve the task and the knowledge available to the methods. Thus, the control of processing is flexible and dynamic rather than rigid or static.

Language Processing

Kurt Eiselt and Ashwin Ram conduct research on natural language processing. Eiselt's research concerns ambiguity resolution and error recovery in natural language understanding. Ram's work deals with story understanding, question asking, and learning in a natural language context. Previous research on language processing in our group has included work on discourse understanding (JUDIS [E. Turner 1989]).

Ambiguity Resolution and Error Recovery

The process of language understanding can be viewed as a series of decisions. For example, consider the reading of a simple story. As each word is read, the reader makes unconscious decisions: Which meaning of this word did the author intend? How does the choice of word meaning affect the meaning of the sentence that is being processed? How does the meaning of this word or sentence affect the overall interpretation of the story? A great deal of research in natural language understanding has been devoted to investigating the mechanisms that allow the language understander to make

these decisions. With few exceptions, these studies have made the overly optimistic assumption that the understander always makes the correct decision.

Because of the ambiguity inherent in natural language, however, it is clear that any understander, whether human or computer, will encounter situations that cause it to make erroneous decisions. As new information is processed, these erroneous decisions can be contradicted and, consequently, must be corrected. Eiselt's recent research focused on the mechanisms that enable the understander to recover from these incorrect decisions with minimal disruption of the understanding process. Eiselt developed a computational model of ambiguity resolution and error recovery during text understanding, called ATLAST, which uses a mechanism for correcting erroneous decisions that is suggested by psycholinguistic studies of human language understanding (Eiselt 1987, 1989). The work on ATLAST not only offers a means of deriving more humanlike performance from natural language interface systems but also suggests constraints on what constitutes a plausible cognitive model of the human language processor.

Question-Driven Understanding

Ram is investigating research issues in the areas of creative understanding, machine learning, and abduction in a natural language-understanding domain. The underlying theme of his work is a focus on the goals of the reasoner and the interaction of these goals with reasoning processes. In particular, his research focuses on *knowledge goals*, the goals of a reasoner to learn by acquiring new knowledge or reorganizing existing knowledge. Knowledge goals, often expressed as questions, arise when the reasoner's model of the domain is inadequate in some reasoning situation. This approach leads the reasoner to focus on what it needs to know, formulate questions to acquire this knowledge, and try to learn by pursuing its questions.

The AQUA project explores these ideas in a natural language-understanding domain. AQUA is a question-driven story-understanding program that learns about terrorism by reading newspaper stories about unusual terrorist incidents in the Middle East

(Ram 1987, 1989). The main point of this research is to create a dynamic story-understanding program that is driven by its questions or goals to acquire knowledge. Rather than being canned, the program changes as its questions change; it reads similar stories differently and forms different interpretations as its questions and interests evolve.

The AQUA project explores issues of learning, explanation, and interestingness in an integrated framework. The intent is not to have the program acquire the right understanding of terrorism but rather to be able to wonder about unusual things it reads about and ask questions about them (Ram 1989). As it learns more about the domain, it asks better and more detailed questions. This kind of questioning forms the origins of creativity; rather than being satisfied with available explanations, a creative person asks questions and tries to explore the explanations in novel ways. One interesting outcome of this work is the formulation of a functional theory of *interestingness* (Ram 1990c): The knowledge goals of the reasoner are used to evaluate the interestingness of various aspects of the stories being read. They also provide a means of controlling the inferences underlying the question-asking process and allow the system to evaluate the interestingness of its questions.

Although the project is being used to explore cognitive issues such as those previously mentioned, there are also practical benefits of a program that can represent and reason about its own goals explicitly. Such a program can focus its limited resources on the relevant aspects of its environment and pay less attention to irrelevant ones. This approach allows it to spend more time drawing those inferences that are relevant and useful to its goals.

New Projects

Several new projects related to natural language processing are under way. In one project, Dorrit Billman (a faculty member in the School of Psychology), Eiselt, and Justin Peterson (a Ph.D. student in the AI group) are investigating the relation between the conceptual structure of events and verb-centered representations of language; they are interested in how this relationship might aid inference in language processing. Many distinctions about event type are syn-

AQUA is a question-driven story-understanding program...

tactically marked in a domain-general model and provide a basis for making inferences about the type of event described. They are interested in identifying how plausible inferences are constrained by particular syntactic distinctions. Because syntactic information does not depend on the particular topic being discussed, it helps the comprehender when faced with *novelty*, the meaning of novel verbs given a familiar use (John skugged the ball to Martha) or the meaning of a familiar verb in a novel use, as in metaphor (The shortstop looked the ball into the glove).

In another project, Goel and Eiselt are using mental models for integrating natural language processing, problem solving, and knowledge acquisition. This research has two components: First, it investigates the use of stored causal models of physical devices, abstract processes, and domain principles to guide the process of natural language interpretation. Second, it explores the acquisition of behavior-structure and causal models of physical devices from their natural language descriptions. An interesting feature of this research is that it uses model-based knowledge representations and reasoning methods that are common to both the language interpreter and the problem solver. These knowledge representations and reasoning methods were used earlier for design problem solving in the KRITIK system.

In a third project, Eiselt and Ram are studying language understanding and commonsense reasoning in the highly complex domain of soap operas, which is characterized by multiple adversarial agents.

A fourth project, led by Ram, is in the area of intelligent information extraction and involves using knowledge goals to focus the understanding process. This project will ultimately lead to two programs, a personalized newspaper program that reads several newspaper stories and puts together a summarized newspaper tailored to the user's interests and

*Scientific data
interpretation can be
modeled as an instance
of abduction...*

an online database program that can retrieve MEDLINE peptide abstracts relevant to the user's interests by using the query to guide the analysis of the abstract. The programs will retrieve information relevant to their initial questions, learn more about the domain by answering these questions, and generate possible questions to pursue further.

Explanation

Kurt Eiselt, Ashok Goel, and Ashwin Ram conduct research on abduction and explanation. Goel's research on abductive explanations is in the domain of scientific data interpretation, Ram's research deals with story understanding, and Eiselt's work concerns language processing. An interesting aspect of Goel and Eiselt's research in explanation is that the abductive architectures they developed have a strong connectionist flavor.

Scientific Data Interpretation

Scientific data interpretation can be modeled as an instance of abduction, that is, as inference to the best explanation for a given set of data. Goel's work on scientific data interpretation viewed as abduction focused on two issues. First, given a set of data to be explained, how can relevant explanatory hypotheses be generated efficiently? Goel formally analyzed several methods for generating explanatory hypotheses (Goel, Soundararajan, and Chandrasekaran 1987; Goel and Bylander 1989). This analysis reaffirms the need for the proper organization of explanatory hypotheses in memory.

Second, given a set of explanatory hypotheses that explain portions of a given data set, how can a best composite explanation be synthesized efficiently? Goel developed a concurrent architecture for abductive reasoning that exploits goal dependencies in synthesizing composite explanations (Goel, Sadayappan, and Joseph-

son 1988). The architecture was successfully simulated on a serial machine and evaluated for medical data interpretation (Fischer and Goel 1990). An especially noteworthy aspect of this abduction architecture is that it can be realized on constraint-optimization neural networks (Goel, Ramanujam, and Sadayappan 1988).

Language Understanding

Eiselt's work on natural language with the ATLAST model concentrated largely on understanding a single sentence or a short text consisting of only a few sentences. To arrive at a coherent interpretation of the text, ATLAST uses a marker-passing search mechanism to find meaningful and, possibly, multiple connections between the words of the input text. These connections represent plausible lexical and pragmatic inferences about the text. ATLAST then applies evaluation metrics to narrow the results of the search to those connections or inferences that best explain the text. This method of abductive inference in ATLAST proved to be well suited to accounting for interesting phenomena in language understanding, such as the online detection and correction of incorrect inferences and individual processing differences in text understanding.

Story Understanding

Ram is investigating the use of case-based reasoning techniques to construct explanations for novel situations encountered by a reasoner. Instead of chaining primitive inference rules, which could be inefficient in complex domains, his program AQUA builds abductive hypotheses from previous cases known to the system (Ram 1990a). Ram is interested in the nature and representation of explanatory cases and the processing issues of constructing and evaluating explanations. An interesting aspect of this work is that the domain knowledge available to the system might be incompletely understood or improperly indexed in memory. This issue is central for abductive inference in complex domains and ties in with Ram's (1990b) work on incremental learning. Ram is also investigating the problem of selecting the best explanation. In the AQUA system, hypotheses are evaluated with respect to the purposes for which they are built in the first place.

Perception

Daryl Lawton and Ronald Arkin conduct research on visual perception. Lawton's research concerns image understanding and several of its applications. Arkin's research deals with action-oriented perception in the context of mobile robotics.

Perceptual Processing

Lawton is interested in all areas of vision research. Of particular importance is perceptual processing for extracting environmental and symbolic information from images for the control of real-time behavior. This work involves three primary areas: motion analysis, perceptual organization, and the incorporation of active sensing strategies into vision. Earlier work in motion analysis (Lawton, Rieger, and Steenstrup 1987) concerned processing restricted cases of motion for which robust solutions are possible (Lawton 1983), real-time motion analysis using a content-addressable parallel processor (Steenstrup, Lawton, and Weems 1983), and techniques for the immediate extraction of motion parameters from the differential properties of optic flow fields at occlusion boundaries (Rieger and Lawton 1985). Current work concerns the psychophysical and practical implications of these approaches and also the direct extraction of occlusion boundaries. It might be unnecessary to use motion analysis to determine an exact depth map or egomotion parameters. A wide range of motor activity and cues for directing attention depend only on extracting and representing the relative depth of surfaces and the image-registered location of occlusion boundaries. Perceptual organization is fundamental for extracting information for model-based recognition and determining landmarks used in qualitative navigation. Current work involves extending the hierarchical rule-based grouper developed in Lawton and McConnell (1987) and Gelband and Lawton (1988).

Model-Based Vision

Model-based vision concerns the use of world knowledge to interpret imagery. Fundamental issues concern such things as how to represent models, preferably as hierarchies of physically and geometrically based constraints; how to index into a potentially enor-

mous database of models during interpretation; and how to perform model-to-image matching, inference, and optimization. Lawton's current work involves building model-based vision systems for several domains, extending the general architecture previously used for outdoor robot landmark extraction and matching (Lawton et al. 1987). Projects under way include one for the analysis of pulmonary embolism (along with researchers in nuclear medicine at Emory Medical School), one for dynamic cardiac images (along with Dr. Norberto Ezquerra of the Medical Informatics Center at Georgia Tech), and one for outdoor robotics. Other efforts are planned for the inspection of chickens and the automatic interpretation of electron density maps obtained from X-ray diffraction imagery.

Many of the difficult issues in model-based vision concern control and hypothesis management. To aid in the development of autonomous systems, Lawton is developing interactive model-based vision systems. These systems have the same underlying architecture as an autonomous system but are controlled by a human. This framework has many exciting implications. The interactive system provides a rich set of protocols for programming the autonomous vision system. Scripts obtained with the interactive system can be used for the transfer of interpretation expertise, especially in areas such as biomedicine. We can also restrict the presentation of imagery to the human to get detailed protocols for psychological study. A current project is developing an interactive vision system for the control of a telerobot where there is a limited communications bandwidth. The human will quickly provide a high-level interpretation of a scene that can be used for the short-term autonomous functioning of a telerobot.

Spatial Understanding

The objective of Lawton's work on spatial understanding is to produce autonomous robots that can freely wander, without harming themselves while they create and improve maps of their environment that can then be used for navigation and planning. Qualitative spatial understanding and navigation, developed with Tod Levitt, concerns spatial learning and path planning in the absence of a

single global coordinate system for describing locations and the positions of landmarks. It is based on a multilevel representation of space, which, at its most abstract level, is based totally on topological properties that allow a robot to describe a location using the directions of visually salient patterns (with no associated range measurements) and then navigate using the occlusions that occur among them as a basic cue to control movement through the environment. An advantage is that the robot can use landmarks for which exact positions or recognition cannot be determined. Thus, if a robot sees a building in the distance, it might not know or be able to recognize the structure as a building or determine its exact position in space, but it can still incorporate this information to form an effective spatial memory for navigation. This approach is actually intuitive: It is doubtful that humans navigate by detecting landmarks, determining ranges to them, and then storing everything in a single frame of reference.

We are currently implementing qualitative navigation on a robot and exploring the necessary perceptual processing (Lawton, Arkin, and Cameron 1990). The spatial representation we developed might have implications for organizing hypermedia-based tutorials as people navigate through overlapping sets of linked information.

Image-Understanding Environments

Over the past 25 years, several different software environments have been developed to support and integrate image-understanding activity. These environments have now evolved into integrated systems of considerable computational and representational power, reflecting the range of problems researchers in computer vision deal with and incorporating much of what has been learned about machine vision. These systems are referred to as image-understanding environments. In addition to previous work in this area (Lawton and McConnell 1988), we are currently working in areas concerned with object-oriented methodologies and facilities for cooperative work for image-understanding environments (Lawton and Mead 1990).

Action-Oriented Perception

Arkin's research on perception focuses on action-oriented perception in the context of mobile robotics. His work is based on the tenet that perception not viewed in the context of motor action is meaningless. This tenet is a departure from the general view of the computer vision community but is consistent with many cognitive psychologists. What this belief affords the robot designer is the ability to produce perceptual algorithms that exploit expectations of what needs to be perceived and that utilize focus-of-attention mechanisms to accomplish this task. Specialized spatial uncertainty management techniques were developed to support this work. The net result is computationally efficient computer vision algorithms that are capable, in the context of specific motor needs, of equal or better performance than more traditional ones (Arkin, Riesenman, and Hanson 1987).

The visual algorithms we developed here at Georgia Tech were geared to support docking operations in a manufacturing environment (Arkin et al. 1989). Temporal activity (motion) detection provides the anticipatory perception required to feed the ballistic component of the docking operator. A spatially constrained Hough transform technique was developed that recognizes the dock at a range of 10 to 20 feet, providing the exteroceptive cue for the transition to controlled motion. Adaptive tracking algorithms then take over. The first uses a fast region segmentation algorithm previously used for road following to recognize a passive landmark placed on the dock. The inverse perspective transform is then used to provide data regarding the relative position of the dock to the robot. A texture-based algorithm then completes the fine positioning of the robot relative to the workstation. We also use ultrasonic data for obstacle avoidance and shaft encoders on the robot to provide coarse information regarding the position of the robot relative to the world.

New Projects

Lawton is developing a generic and extendable sensor platform for a mobile robot consisting of inexpensive off-the-shelf items. These items include multiple computer-con-

trolled cameras, an inclinometer, and accelerometers and will be attachable through ethernet to several different workstations.

Arkin and Erika Rogers (a Ph.D. student with the AI group) are currently working with human subjects in the context of biomedical image processing with the intent of developing a cognitive model of human perception in this domain. The immediate research goal is to use this derived model, based on considerable experimentation with human subjects, to provide a predictive model for image-enhancement techniques. From a pragmatic viewpoint, this model will give rise to more efficient processing by radiologists in terms of diagnostic accuracy and throughput. Although considerable research has been conducted in the area of radiological image interpretation, we have not found any examples where image-enhancement techniques were derived from an understanding or model of the perceptual and cognitive processing in humans. Our approach is unique in this respect. More often, an image-processing technique is applied to an image, and

then the performance of the radiologists is analyzed. The technique chosen generally does not have any theoretical basis for its use. In our work, by first studying the perceptual and diagnostic performance of radiologists in their natural environment and then slowly and carefully altering the parameters and dimensions of their task, we hope to produce a coherent model that can result in marked benefits for this problem domain as well as fodder for techniques that can be applied to high-level cognitive robotic perception. This research can potentially provide insights into learning because one dimension of the study involves the differences between expert and non-expert radiologists. This research is being conducted jointly with the Emory University School of Medicine.

Robotics

Ronald Arkin also conducts research in robotics. His research in this area is concerned with intelligent navigation and reactive control.

Robot Navigation

Arkin's research provides a coherent framework for uniting the beneficial aspects of reactive control and hierarchical planning. Reactive control is intimately tied to sensory feedback, providing the ability to react to unanticipated events in an efficient manner. When serving as the sole basis for navigation, purely reactive systems suffer many pitfalls and generally can only give rise to intelligence at the level of animals. His system (AURA, the autonomous robot architecture) incorporates many of the beneficial aspects of hierarchical control. The uniqueness of this approach lies in the ability to dynamically reconfigure motor and perceptual behavior networks to match the current environment and mission needs of the mobile vehicle (Arkin 1990).

Additionally, his approach is unique in its formulation of motor behaviors and perceptual strategies for mobile robot navigation as schemas (Arkin 1989b). These representation and control units correlate strongly with models of human and

animal behavior. Although it is not his goal to build robots that strictly mimic humans, he approaches the problem realizing that we can benefit by studying theories of biological systems that already can solve the difficult problem of navigation in a changing world.

It has always been a goal of his research to derive general-purpose navigational techniques rather than design ad hoc solutions. Navigational path-planning techniques have been applied to a variety of new domains, including three-dimensional navigation that could be found in space and undersea robotics. Extensions of the world representations were made that permit navigational path planning for a flying-crawling robot. We also developed schema-based control techniques for free-swimming or free-flying robots, which were successfully tested in simulation (Arkin 1989c).

We also studied navigation in undulating terrain (Arkin and Gardner 1990). This research is particularly useful for surveillance operations. These techniques enable us to literally implement AI hill-climbing behavior as well as isocontour following and valley seeking. His use of a stochastic background process to deal with local minima and maxima is unique.

As described in the previous section, we are also studying the application of mobile robotics to manufacturing, with the goal of replacing automatic guided vehicles with more intelligent machines. His coupling of world modeling with reactive control is distinct from other approaches to this problem (Arkin and Murphy 1990).

Robot Survivability

For robots to be truly autonomous, we must study how to make them survive in situations that threaten their existence. Reactive control techniques, as previously described, enable a robot to cope with a rapidly changing world. Arkin is also looking at other survival techniques that are more introspective: homeostatic control and fault-tolerant control systems.

Homeostatic Control Arkin's approach to this problem is unique, and we are one of the few research centers studying it. *Homeostatic control* involves dynamic replanning in hazardous environments (Arkin 1989a). This technique is fully integrated in concept and simulation with the reactive control system we

Several members of our group conduct research in the representation and organization of knowledge.

developed within Aura. The intent is to produce smooth, graceful alterations in motor behavior based on changing internal conditions, not only external environmental perceptions. We demonstrated in simulation the importance and feasibility of this approach. An analog of the mammalian endocrine system is forwarded as the best means for supporting this type of control. Broadcast communication mechanisms with signal schemas embedded at receptor sites within motor schemas enable the simulation of feast-or-famine behavior as well as more gradual behavior variations based on available fuel and internal temperatures for the vehicle.

Fault-Tolerant Control Another area of robot survivability concerns what to do when hardware failure occurs. Arkin is studying this problem jointly with faculty in electrical engineering (Arkin and Vachtsevanos 1990). When hardware component failure occurs in a complex engineered system, the fault usually propagates through the system. This research centers on the issues of fault detection, isolation, propagation, and reconfiguration. The goal is to enable a system to reconfigure itself in real time by using qualitative models of the system to approximate the extent of the fault. After these AI techniques are applied, the affected subsystems can be reconfigured using more traditional control techniques. This approach enables the system to operate in a suboptimal yet survivable mode, buying time until it can either be repaired or fully reconfigured. The importance of this work for robotic operations in space, rescue operations, and hazardous environments is self-evident.

New Projects

A new project on cognitively based models for sensor fusion is being investigated by Arkin and Robin

Murphy (a Ph.D. student in the AI group). This novel approach exploits perceptual information inherent in the environment, extracting only the perceptual data necessary to support a given set of actions. General mechanisms are being designed that are independent of contributing sensor modalities, providing a bias toward a dominant source when appropriate. This work is strongly motivated by cognitive psychology and is intended to provide the advantages of a more tractable fusion methodology with improved global results.

Knowledge Representation and Organization

Several members of our group conduct research in the representation and organization of knowledge.

Janet Kolodner (1983a, 1983b), for example, has been investigating the representation and organization of experiences in memory for more than a decade. Ronald Arkin's research concerns the issues of representational adequacy in the context of schema-based navigation. Ashok Goel conducts research on the representation and organization of mental models of teleological artifacts. Ashwin Ram is interested in the representation and organization of explanatory cases for motivational analysis and abduction.

Models

Goel is investigating behavior-structure models of teleological artifacts, such as physical devices. A behavior-structure model of a physical device represents knowledge of how the output behaviors of the artifact, which include its functions, arise from its structure. The function of a flashlight, for example, is to create light when a switch is pressed. Its structure consists of a battery, a switch, and a bulb connected in series. The behavior-structure model of the flashlight represents knowledge of how the structure of the flashlight results in its output behaviors, including its function of creating light when the switch is pressed.

Goel (1989) developed a specific behavior-structure model, the behavioral component-substance model, for a class of physical devices. This model is applicable to devices whose functioning can be viewed in terms of flow of substances from one com-

ponent to another, including abstract substances such as heat. The behavioral component-substance model explicitly specifies the expected output behaviors of the design, including its functions, the elementary structural and behavioral interactions between the components and substances constituting the structure of the design, and the internal causal behaviors of the design that compose the elementary interactions into its output behaviors. The output behaviors in this model are represented as schemas. The schema for an output behavior specifies the causal behavior responsible for it. These causal behaviors are represented as directed acyclic graphs. The states and the transitions in the causal behaviors are themselves represented as schemas. The schema corresponding to a transition specifies the causes underlying the transition, where the causes are specified in terms of structural and behavioral interactions between structural components, physics principles, qualitative equations, and so on. This behavioral component-substance model was instantiated in the KRITIK system and evaluated for use in adaptive design problem solving.

Explanations

Ram is investigating the nature and representation of volitional explanations for motivational analysis of natural language stories. Volitional explanations are constructed from *decision models*, which describe the planning process that an agent goes through when considering whether to perform an action. Decision models are represented as *explanation patterns*, which are standard patterns of causality based on previous experiences of the understander. These patterns are indexed in memory using anomaly, situation, and stereotype indexes. Ram (1990a) is interested in the nature of explanation patterns; their use in representing decision models; and the process by which they are retrieved, used, and evaluated.

New Projects

Goel is exploring several new lines of research on the representation and organization of mental models of physical devices. He is extending the behavioral component-substance model to include the representation of fields such as the magnetic field.

Although substances can flow from one component to another only if the two components are spatially connected, fields can influence at a distance and, therefore, require a different representation. He is also investigating the organization of mental models in memory. He is developing an indexing scheme in which the models are indexed using functional, structural, and process descriptions. Finally, he is investigating the representation and organization of device-independent abstract models of generic mechanisms, such as feedback, and domain-independent abstract models of physics principles and equations, such as the laws of thermodynamics.

Learning and Instruction

Several members of our group conduct research on the acquisition of knowledge, learning, and instruction. Janet Kolodner, for example, conducts research in case-based learning; she is also investigating the learning of plans. Ashwin Ram is investigating explanation-based and case-based learning. Ashok Goel is exploring model-based learning and explanation-based learning.

Learning Plans

In the EXPEDITOR project, Kolodner and Steve Robinson (a Ph.D. student in the AI group) are investigating the learning of plans. EXPEDITOR is responsible for scheduling the tasks and subtasks of a single parent. It begins by knowing how to do individual tasks (for example, get the kids dressed, eat breakfast, give the kids breakfast, find a babysitter, do laundry). Over time, it learns two things: how to interleave these tasks in appropriate ways to achieve its goals efficiently and how to execute these tasks in the context of anticipated and recurring failures.

Thus, it learns, for example, how to interleave getting dressed, getting the kids dressed, eating breakfast, giving the kids breakfast, and so on, in such a way that it can get out of the house in the morning in reasonable time. In addition, it learns that when there is no breakfast cereal, peanut butter and jelly sandwiches make a fine breakfast.

EXPEDITOR can be viewed as the prototype for logistics scheduling in domains (1) where much interleaving of plans must be done, but the

same interleaving can be repeated over and over and (2) where problem situations happen often, but the same problem situations recur and can be anticipated. MEDIC's goal scheduler supplies some of the control needed for such systems but does not address the learning component. Learning about plan interleaving, problem anticipation, and fixes that work are EXPEDITOR's main tasks. Goal scheduling and goal tracking are two important problems that must be addressed for learning to happen. Kolodner is looking for a goal-tracking mechanism that can deal with the huge numbers of goals present in an open world but still allow the efficient execution of well-known plans. Currently, EXPEDITOR's efficiency comes from its case representations and dynamic memory organization (Kolodner 1983a, 1983b).

Case-Based Learning

In the CELIA project, Kolodner and Michael Redmond (who is writing his Ph.D. thesis on this project) are investigating case-based learning. Like MEDIC, CELIA addresses diagnosis as a planning problem. To diagnose a car, one must first verify the customer's complaint, then come up with a hypothesis of what is wrong with it, refine the hypothesis to something that can be tested, test it, interpret the test, and so on. CELIA uses cases to both guide its diagnostic planning (what subgoal to pursue next) and suggest solutions (for example, the problem is in the carburetor). Two case-based reasoning problems are addressed in CELIA: how to represent and index cases so that full cases or case pieces can be accessed (Redmond 1990) and how to acquire cases from a teacher (that is, how to turn a teacher's examples into cases that can be stored in the student's memory for later use) (Redmond 1989b).

In CELIA, Kolodner and Redmond are investigating the role that concrete problems play in learning a problem-solving task. They are studying the processes by which a student diagnostician can learn by solving new problems that are just beyond its abilities and explaining the actions and commentary of a teacher. CELIA, as previously described, is a computer program that diagnoses automobile engine problems. In addition to diagnosing problems, CELIA watches a teacher solve diag-

nostic problems and learns from the teacher (Redmond 1989a). As a result, its own diagnostic skills improve. In learning from the teacher, CELIA uses its diagnostic knowledge to predict what the teacher will do next. In cases where its predictions don't match what the teacher is doing, CELIA explains the anomaly. Based on the teacher's explanations, CELIA sometimes fills gaps in its knowledge and sometimes refines what it already knows. As CELIA watches the teacher, it creates cases from the teacher's examples and saves these cases as its own. The cases give it a way to index into the knowledge it has of car mechanics.

CELIA is based on observations we made of student car mechanics at several different levels of skill (Lancaster and Kolodner 1987). Kolodner saw students using several different learning strategies. Active gap filling, learning by understanding explanations (Martin and Redmond 1988), and case-based reasoning were among the strategies we observed. CELIA implements these three strategies. We learned several important things from this project: First is the active role students must play in learning from a teacher. Merely listening is not enough. The prediction of what the teacher will do and the comparison of these predictions to what is actually done by the teacher are necessary for the student to figure out what it is s/he is supposed to learn. Second is the importance of a teacher's explanations in novice skill learning. Many things can directly be learned from experience, but some things are hidden and require a teacher's explanation. This situation might be because of the bias of the learner; it might be because of missing knowledge; or it might be because a long causal chain must be constructed to connect two pieces of knowledge, longer than resources usually allow.

Explanation-Based Learning and Case-Based Learning

Ram is interested in explanation-based and case-based learning. His research investigates how a reasoner can improve its understanding of an incompletely understood domain by applying what it already knows to novel problems in this domain. The performance of a case-based reasoning system depends on its case library of past experiences. In a novel

and complex domain, however, situations encountered previously by the reasoner might not have been completely understood. Furthermore, the reasoner might not even have a case that adequately deals with the new situation or might not be able to access the case using existing indexes.

Ram developed a model of incremental case-based learning that allows a reasoner to (1) use domain knowledge that might not be completely understood to solve novel problems, (2) maintain an explicit model of the gaps in its knowledge base, (3) learn by filling in these gaps when the information it needs becomes available, and (4) gradually evolve a better understanding of the domain (Ram 1989, 1990c). The learning process is focused by the knowledge goals of the reasoner to acquire and organize domain knowledge. Additional constraints are derived from the causality of the domain using explanation-based learning techniques. Ram's program incrementally improves its domain model by modifying known but incompletely understood explanation patterns as well as learning new indexes for explanation patterns. Ram is also investigating methods for learning through explanation-based refinement or specialization of abstract knowledge.

In conjunction with the theory of learning, Ram (1989) is also investigating the nature of knowledge goals. He is developing a theory of the types of knowledge goals, an opportunistic memory architecture for the management of knowledge goals, and a theory of explanation that characterizes the kinds of explanations that constitute satisfactory solutions to the problems that give rise to knowledge goals.

Model-Based Learning and Case-Based Learning

Goel's research on model-based learning and case-based learning focuses on two issues. First, he is investigating the learning of behavior-structure models of teleological artifacts. He developed a computational theory for incremental learning of behavior-structure models of physical devices. In this scheme, new behavior-structure models are learned by adapting and revising old ones (Goel 1989). That is, the behavior-structure model of one physical device is learned by revising the

model for a similar device. The process of model revision is focused by the specific differences between the structures of the two devices. This scheme for incremental learning of models was successfully tested in the KRITIK system. Because the behavior-structure model of a teleological artifact provides a causal explanation of how the structure of the artifact results in its output behaviors, model-based learning in this context is closely related to explanation-based learning.

Second, Goel is investigating the use of mental models in solving a class of credit-assignment problems. The proper assignment of credit is a major issue in design adaptation. For example, given the functional specifications of a desired design and the functional and structural specifications of a known design, the credit-assignment problem is to decide the structural causes for the failure of the known design to achieve the desired functions. In his work on the KRITIK system, Goel showed how the behavior-structure model for the known design helps to localize the structural causes and identify the needed structural modifications. Once these structural modifications are verified and executed, the new design can be stored in memory as another design case.

New Projects

Several new research projects on knowledge acquisition, learning, and instruction are under way. Kolodner is directing two new projects: example-based instruction and the design of lesson plans. In the project on example-based instruction, Kolodner is looking at the implications for teaching what she found out about learning. She is interested in the kinds of teaching strategies that facilitate learning and the kinds of instruction in learning that should be given to students to help them to learn. More specifically, she is investigating the role of examples in teaching and the design of systems that can store examples in an example library and use them for such things as motivating new ideas, making abstract principles concrete, inducing generalizations, and inducing discriminations. Her intention is to build an example library that, in response to teacher goals, will create examples the teacher can present to students.

You are cordially invited to become a member of the AI Community's principal technical society:

The American Association for Artificial Intelligence

CONSIDER THESE BENEFITS...

- **AI Magazine** — A quarterly publication (called the "Journal of Record for the AI Community" by *High Technology* magazine) devoted to the entire field of artificial intelligence.
- **AAAI Membership Directory** — An annual roster of AAAI members sharing an interest in artificial intelligence.
- **Proceedings of the National Conference on Artificial Intelligence** — An annual publication consisting of state of the art research papers presented at the AAAI's annual conference. (This publication is included in conference registration; AAAI members not attending the conference may purchase the Proceedings at a reduced rate.)
- **Technical Program of the National Conference on Artificial Intelligence** — Highlighting examples of significant contributions in the development of state of the art in AI practice. (AAAI members receive a substantially reduced rate on conference registration fees.)
- **Exhibit Program of the National Conference on AI** — Affording attendees the opportunity to view the latest in AI products, services, and research.
- **Tutorial Series** — On state-of-the-art topics. (AAAI members receive a reduced rate on tutorial registration fees.)
- **Substantially reduced subscription rate for *Artificial Intelligence***
- **Assistance in formation of local and regional AI groups**

To take advantage of all these benefits, complete and send the coupon together with your check (in U.S. funds only) for \$40.00 (\$65.00 foreign) to the address below (or call 415/328-3123 for further details). Send to:

AAAI Membership • 445 Burgess Drive • Menlo Park, California 94025

Name	_____
Company/Mail Stop	_____
Address	_____
City	_____
State/Zip/Country	_____

In the project on designing lesson plans, Kolodner's aim is to create an interactive computer system that can help elementary school teachers create lesson plans for teaching science. This project combines our interests in design, planning, and instruction in an area where there is a real need for innovation. Her intention is to help teachers create a flexible plan that includes motivating examples appropriate to the experiences of the students, guidelines for explaining and reexplaining difficult points, pointers on what to emphasize, guidelines that allow the teacher to notice problems students are having and areas of extra interest among the students, and guidelines for addressing these special areas. Teachers are treated as *opportunistic planners*, that is, planners that can flexibly change their plans if the need or opportunity arises. The lesson plan that is created will be a sort of contingency plan, and teachers will be made aware of circumstances that require following alternate branches of the plan and opportunities that might arise for adding something appropriate to the plan.

Ram is leading several new projects on learning, including incremental learning of explanation patterns and learning new indexes for cases in memory. He is extending his previous ideas to incremental learning by case-based reasoning across domains rather than only within the same domain. He is also interested in the issue of knowledge planning and the development of a system that learns by planning to satisfy its knowledge goals.

Goel is directing two new projects on model-based learning. In one project, he is extending his previous work on model revision to physical devices whose internal causal behaviors can only partially be ordered. In another project, he is exploring the role of device-independent process models in revising device-specific behavior-structure models.

Beyond the AI Group in the College of Computing

Beyond the AI group, several other research groups at Georgia Tech are engaged in active research in areas

related to AI. Within the College of Computing, several faculty members conduct research on AI-related areas such as human-computer interaction, databases and knowledge bases, and neural networks. The School of Psychology supports a strong group in the area of cognitive psychology. This group currently consists of four faculty members, Lawrence Barsalou, Dorrit Billman, Susan Boviar, and Richard Catrambone, all of whom actively collaborate with the AI group. In one collaborative project, for example, Billman, Joel Martin (who is writing his Ph.D. dissertation on this project), and Kolodner are investigating the learning of predictive knowledge from examples (Martin 1989; Martin and Billman 1991).

Similarly, the Man-Machine Systems Research Center in the School of Industrial and Systems Engineering has a strong research program in cognitive engineering. The three faculty members in this group, T. Govindraj, Alexander Kirlik, and Christine Mitchell, also interact with the AI group. Several other engineering departments also conduct research on

some aspects of applied AI, including expert systems, cognitive engineering, and robotics. The Georgia Tech Research Institute has a large research and development program in several aspects of applied AI, including expert systems, cognitive engineering, robotics, and neural networks. In the School of Language, Communication, and Culture, Gregory Colomb is investigating language processing from a cognitive perspective.

The AI group actively cooperates with many of these groups, especially the cognitive psychology group in the School of Psychology. Recently, our former School of Information and Computer Science was elevated to a full College of Computing, in part to facilitate interdisciplinary interaction and collaborative research. The college, with a new dean, is housed in a new building. This year we are starting a cognitive science Ph.D. certificate program in cooperation with our colleagues in other academic units. A research center for cognitive science is being planned. The future of AI at Georgia Tech looks bright indeed!

Acknowledgments

AI research in the College of Computing is supported by grants from the National Science Foundation, the National Aeronautics and Space Administration, the National Institute of Health, the Defense Advanced Research Projects Agency, the Office of Naval Research, the Army Research Institute, the Army Engineering Topographic Laboratory, Advanced Decision Systems, International Business Machines, Emory/Georgia Tech Biomedical Research Center, Georgia Tech Research Foundation, Georgia Tech's Computer-Integrated Manufacturing Systems Program, and the College of Computing at Georgia Tech. Jeannie Terrel helped with the preparation of this document.

Note

1. See the article by Janet Kolodner in this issue.

Bibliography

Arkin, R. 1990. Integrating Behavioral, Perceptual, and World Knowledge in Reactive Navigation. *Robotics and Autonomous Systems* 6(1): 105–122.

Arkin, R. 1989a. Dynamic Replanning for a Mobile Robot Based on Internal Sensing. In Proceedings of the 1989 IEEE Interna-

tional Conference on Robotics and Automation, 1416–1421. Washington, D.C.: IEEE Computer Society.

Arkin, R. 1989b. Motor Schema-Based Mobile Robot Navigation. *International Journal of Robotics Research* 8(4): 92–112.

Arkin, R. 1989c. Three-Dimensional Motor Schema-Based Navigation. In Proceedings of the NASA Conference on Space Telerobotics, 291–300. Pasadena, Calif.: Jet Propulsion Laboratory.

Arkin, R., and Gardner, W. 1990. Reactive Inclinometer-Based Mobile Robot Navigation. In Proceedings of the IEEE Conference on Robotics and Automation, 936–941. Washington, D.C.: IEEE Computer Society.

Arkin, R., and Murphy, R. 1990. Autonomous Navigation in a Manufacturing Environment. *IEEE Transactions on Robotics and Automation* 6(4): 445–454.

Arkin, R., and Vachtsevanos, G. 1990. Qualitative Fault Propagation in Complex Systems. In Proceedings of the Twenty-Ninth IEEE Conference on Decision and Control, 1509–1510. Washington, D.C.: IEEE Computer Society.

Arkin, R.; Riseman, E.; and Hanson, A. 1987. Visual Strategies for Mobile Robot Navigation. In Proceedings of the IEEE Computer Society Workshop on Computer Vision, 176–181. Washington, D.C.: IEEE Computer Society.

Arkin, R.; Murphy, R.; Pearson, M.; and Vaughn, D. 1989. Mobile Robot Docking Operations in a Manufacturing Environment: Progress in Visual Perceptual Strategies. In Proceedings of the IEEE International Workshop on Intelligent Robots and Systems, 147–154. Washington, D.C.: IEEE Computer Society.

Eiselt, K. 1989. Inference Processing and Error Recovery in Sentence Understanding. Ph.D. diss., Dept. of Information and Computer Science, Univ. of California at Irvine.

Eiselt, K. 1987. Recovering from Erroneous Inferences. In Proceedings of the Sixth National Conference on Artificial Intelligence, 540–544. Menlo Park, Calif.: American Association for Artificial Intelligence.

Fischer, O., and Goel, A. 1990. Abductive Explanation: On Why the Essentials Are Essential. In *Intelligent Systems Methodologies*, volume 5, eds. Z. Ras, M. Zemenkova, and M. Emrich, 354–361. Amsterdam: North Holland.

Gelband, P., and Lawton, D. 1988. Perceptual Grouping Network Architecture. In Proceedings of the Conference of the International Neural Network Society), 802–803. Washington, D.C.: IEEE Computer Society.

Goel, A. 1990. Grounding Case Adaptation in Causal Models. In *Intelligent Systems Methodologies*, volume 5, eds. Z. Ras, M. Zemenkova, and M. Emrich, 260–267. Amsterdam: North Holland.

Goel, A. 1989. Integration of Case-Based Reasoning and Model-Based Reasoning for Adaptive Design Problem Solving. Ph.D. diss., Dept. of Computer and Information Science, The Ohio State Univ.

Goel, A., and Bylander, T. 1989. Computational Feasibility of Structured Matching. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 11(12): 1312–1316.

Goel, A.; Callantine, T.; Shankar, M.; and Chandrasekaran, B. 1991. Representation, Organization, and the Use of Topographic Models of Physical Spaces for Route Planning. In Proceedings of the Seventh Conference on Artificial Intelligence Applications, 308–314. Washington, D.C.: IEEE Computer Society.

Goel, A., and Chandrasekaran, B. 1990. A Task Structure for Case-Based Design. In Proceedings of the IEEE International Conference on Systems, Man, and Cybernetics, 587–592. Washington, D.C.: IEEE Computer Society.

Goel, A., and Chandrasekaran, B. 1989a. Functional Representation of Designs and Redesign Problem Solving. In Proceedings of the Eleventh International Joint Conference on Artificial Intelligence, 1388–1394. Menlo Park, Calif.: International Joint Conferences on Artificial Intelligence.

Goel, A., and Chandrasekaran, B. 1989b. Use of Device Models in Design Adaptation. In *Proceedings of the Second DARPA Workshop on Case-Based Reasoning*, 100–109. San Mateo, Calif.: Morgan Kaufmann.

Goel, A.; Ramanujam, J.; and Sadayappan, P. 1988. Toward a Neural Architecture for Abductive Reasoning. In Proceedings of the Second IEEE International Conference on Neural Networks, volume 2, 681–688. Washington, D.C.: IEEE Computer Society.

Goel, A.; Sadayappan, P.; and Josephson, J. 1988. Concurrent Synthesis of Composite Explanatory Hypotheses. In Proceedings of the Seventeenth International Conference on Parallel Processing, volume 3, 156–160. University Park, Penn.: Pennsylvania State University Press.

Goel, A.; Soundararajan, N.; and Chandrasekaran, B. 1987. Complexity in Classificatory Reasoning. In Proceedings of the Sixth National Conference on Artificial Intelligence, 421–425. Menlo Park, Calif.: American Association for Artificial Intelligence.

Hinrichs, T. 1990. Some Criteria for Evaluating Designs. In *Proceedings of the Twelfth Annual Conference of the Cognitive Science Society*, 900–907. Hillsdale, N.J.: Lawrence Erlbaum.

Hinrichs, T. 1989. Strategies for Adaptation and Recovery in a Design Problem Solver. In *Proceedings of the 1989 DARPA Workshop on Case-Based Reasoning*, 115–118. San Mateo, Calif.: Morgan Kaufmann.

Hinrichs, T. 1988. Toward an Architecture for Open World Problem Solving. In *Proceedings of the 1988 DARPA Workshop on*

- Case-Based Reasoning*, 182–189. San Mateo, Calif.: Morgan Kaufmann.
- Kolodner, J. 1990. An Introduction to Case-Based Reasoning, Technical Report, GIT-ICS-90-19, School of Information and Computer Science, Georgia Institute of Technology.
- Kolodner, J. 1989. Retrieving the Best Case for a Case-Based Reasoner. In *Proceedings of the Eleventh Annual Conference of the Cognitive Science Society*, 155–162. Hillsdale, N.J.: Lawrence Erlbaum.
- Kolodner, J. 1987. Extending Problem-Solving Capabilities through Case-Based Inference. In *Proceedings of the Fourth Machine Learning Workshop*, 167–178. San Mateo, Calif.: Morgan Kaufmann.
- Kolodner, J. 1983a. Maintaining Organization in a Conceptual Memory for Events. *Cognitive Science* 7: 243–280.
- Kolodner, J. 1983b. Reconstructive Memory: A Computer Model. *Cognitive Science* 7:243–280.
- Kolodner, J., and Kolodner, R. 1987. Using Experience in Clinical Problem Solving: Introduction and Framework. *IEEE Transactions on Systems, Man, and Cybernetics* 17:420–431.
- Kolodner, J., and Penberthy, L. 1990. A Case-Based Approach to Creative Problem Solving. In *Proceedings of the Twelfth Annual Conference of the Cognitive Science Society*, 978–985. Hillsdale, N.J.: Lawrence Erlbaum.
- Kolodner, J., and Simpson, R. 1989. The MEDIATOR: Analysis of an Early Case-Based Reasoner. *Cognitive Science* 13:507–550.
- Kolodner, J., and Simpson, R. 1984. Experience and Problem Solving: A Framework. In *Proceedings of the Sixth Conference of the Cognitive Science Society*, 239–243. Hillsdale, N.J.: Lawrence Erlbaum.
- Kolodner, J.; Simpson, R.; and Sycara, K. 1985. A Process Model of Case-Based Reasoning in Problem Solving. In Proceedings of the Ninth International Joint Conference on Artificial Intelligence, 284–290. Menlo Park, Calif.: International Joint Conferences on Artificial Intelligence.
- Lancaster, J., and Kolodner, J. 1987. Problem Solving in a Natural Task Domain as a Function of Experience. In *Proceedings of the Ninth Annual Conference of the Cognitive Science Society*, 727–736. Hillsdale, N.J.: Lawrence Erlbaum.
- Lawton, D. 1983. Processing Translational Motion Sequences. *Computer Vision, Graphics, and Image Processing* 22:116–144.
- Lawton, D., and McConnell, C. 1988. Image Understanding Software Environments. In Proceedings of the IEEE (Special Issue on Computer Vision, 1036-1049). Washington, D.C.: IEEE Computer Society.
- Lawton, D., and McConnell, C. 1987. Perceptual Organization Using Interestingness. Presented at the AAAI Workshop on Spatial Reasoning and Multisensor Fusion, 13–17 July, Seattle, Wash.
- Lawton, D., and Mead, D. 1990. A Modular Object-Oriented Image Understanding Environment. In Proceedings of the IEEE International Conference on Pattern Recognition, 864–871. Washington, D.C.: IEEE Computer Society.
- Lawton, R.; Arkin, R.; and Cameron, J. 1990. Qualitative Spatial Understanding and Reactive Control for Autonomous Robots. In Proceedings of the IEEE International Workshop on Intelligent Robotic Systems. Washington, D.C.: IEEE Computer Society.
- Lawton, D.; Rieger, J.; and Steenstrup, M. 1987. Computational Techniques in Motion Processing. In *Vision, Brain, and Cooperative Computation*, eds. M. Arbib and A. Hanson, 419–488. Cambridge, Mass.: MIT Press.
- Lawton, T.; Levitt, C.; McConnell, C.; Nelson, P.; and Glicksman, J. 1987. Terrain Models for an Autonomous Land Vehicle. In Proceedings of the IEEE International Conference on Robotics, 2043–2051. Washington, D.C.: IEEE Computer Society.
- Levitt, T. and Lawton, D. 1990. Qualitative Navigation for Mobile Robots. *Artificial Intelligence* 44(3): 305–360.
- Martin, J. 1989. Reducing Redundant Learning. In *Proceedings of the Sixth International Workshop on Machine Learning*. San Mateo, Calif.: Morgan Kaufmann.
- Martin, J., and Billman, D. 1991. Representational Specificity and Learning. In *Computational Approaches to Concept Formation*, eds. D. Fisher and P. Langley. San Mateo, Calif.: Morgan Kaufmann. Forthcoming.
- Martin, J., and Redmond, M. 1988. The Use of Explanations for Completing and Correcting Causal Models. In *Proceedings of the Tenth Annual Conference of the Cognitive Science Society*, 440–446. Hillsdale, N.J.: Lawrence Erlbaum.
- Ram, A. 1990a. Decision Models: A Theory of Volitional Explanation. In *Proceedings of the Twelfth Annual Conference of the Cognitive Science Society*, 198–205. Hillsdale, N.J.: Lawrence Erlbaum.
- Ram, A. 1990b. Incremental Learning of Explanation Patterns and Their Indices. In *Proceedings of the Seventh International Conference on Machine Learning*, 313–320. San Mateo, Calif.: Morgan Kaufmann.
- Ram, A. 1990c. Knowledge Goals: A Theory of Interestingness. In *Proceedings of the Twelfth Annual Conference of the Cognitive Science Society*, 206–214. San Mateo, Calif.: Morgan Kaufmann.
- Ram, A. 1989. Question-Driven Understanding: An Integrated Theory of Story Understanding, Memory, and Learning. Ph.D. diss., Dept. of Computer Science, Yale Univ.
- Ram, A. 1987. AQUA: Asking Questions and Understanding Answers. In Proceedings of the Sixth National Conference on Artificial Intelligence, 312–316. Menlo Park, Calif.: American Association for Artificial Intelligence.
- Redmond, M. 1990. Distributed Cases for Case-Based Reasoning; Facilitating Use of Multiple Cases. In Proceedings of the Ninth National Conference on Artificial Intelligence, 304–309. Menlo Park, Calif.: American Association for Artificial Intelligence.
- Redmond, M. 1989a. Combining Explanation Types for Learning by Understanding Instructional Examples. In *Proceedings of the Eleventh Annual Conference of the Cognitive Science Society*, 147–154. Hillsdale, N.J.: Lawrence Erlbaum.
- Redmond, M. 1989b. Learning from Others' Experience: Creating Cases from Examples. In *Proceedings of the Second Workshop on Case-Based Reasoning*, 309–312. San Mateo, Calif.: Morgan Kaufmann.
- Rieger, J., and Lawton, D. 1985. Processing Differential Image Motion. *Journal of the Optical Society of America* 2:354.
- Rogers, E.; Arkin, R.; Baron, M.; Ezquerro, N.; and Garcia, E. 1990. Visual Protocol Collection for the Enhancement of the Radiological Diagnosis Process. In Proceedings of the First Conference on Visualization in Biomedical Computing, 208–215. Washington, D.C.: IEEE Computer Society.
- Shinn, H. 1989. A Unified Approach to Analogical Reasoning. Ph.D. diss., School of Information and Computer Science, Georgia Institute of Technology.
- Steenstrup, M.; Lawton, D.; and Weems 1983. Determining the Rotational and Translational Components of Optic Flow Fields Using a Content-Addressable Parallel Processor. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, 401–404. Washington, D.C.: IEEE Computer Society.
- Sycara, K. 1987. Resolving Adversarial Conflicts: An Approach Integrating Case-Based and Analytic Methods. Ph.D. diss., School of Information and Computer Science, Georgia Institute of Technology.
- Sylvan, D.; Goel, A.; and Chandrasekaran, B. 1990. Analyzing Political Decision Making from an Information Processing Perspective: JESSE. *American Journal of Political Science* 34(1): 74–123.
- Turner, E. 1989. Interpreting Intention and Convention to Organize Problem-Solving Dialogues. Ph.D. diss., School of Information and Computer Science, Georgia Institute of Technology.
- Turner, R. 1989. A Schema-Based Model of Adaptive Problem Solving. Ph.D. diss., School of Information and Computer Science, Georgia Institute of Technology.

Ashok K. Goel received a Master's in physics in 1980, a Master's in computer and information science in 1982, and a

Applied AI News

Continued from page 16

Carnegie Group (Pittsburgh, PA) has been awarded a \$225,000 grant from the National Science Foundation to continue its study of developing a hybrid neural network/expert system technology. The research will address such issues as comparing various neural networks for effective signal classification and exploring techniques for explaining the hybrid system's reasoning to the system developer and end-user.

GardenTech (Redondo Beach, CA) has developed Bugs, an expert system for gardeners that enables users to identify insect pests and determine which are damaging vegetables and orchard crops. The system recommends environmentally safe pest control, rather than pesticide use.

Avanti Systems International (Batesville, AR) has developed a voice-activated dental charting system. The Victor Voice Chart System allows a dental hygienist to conduct a periodontal examination and dictate findings to an intelligent assistant system.

This news supplement is a service of *ISR: Intelligent Systems Report*, the newsletter of applied artificial intelligence. For further details on the above stories, or to report information on other applications of AI, contact ISR at 2555 Cumberland Parkway, Suite 299, Atlanta, GA 30339, (404) 434-2187, FAX: (404) 432-6969. Subscription price: \$249 for 12 issues.



Ph.D. in computer and information science in 1989, all from The Ohio State University. He is currently an assistant professor with the College of Computing at the Georgia Institute of Technology. He has published a number of papers on planning and design problem solving, abductive inference and explanation, knowledge representation and memory organization, and knowledge acquisition and learning. He chairs the Artificial Intelligence Technical Committee of the Institute of Electrical and Electronic Engineers Systems, Man, and Cybernetics Society.



Ronald C. Arkin received a bachelor of science from the University of Michigan in 1971, a Master's in science from Stevens Institute of Technology in 1977, and a Ph.D. in computer science from the University of Massachusetts at Amherst in 1987. From 1977 until 1985, he was a member of the faculty of Hawthorne College in Antrim, New Hampshire. After receiving his Ph.D., he assumed the position of assistant professor in the College of Computing at the Georgia Institute of Technology, where he was promoted to associate professor in 1990.



Kurt P. Eiselt is an assistant professor in the College of Computing at the Georgia Institute of Technology. He received a Ph.D. from the University of California at Irvine for his work in natural language understanding. His current research interests include the effects of working memory constraints on text understanding, reasoning in highly complex domains, and the acquisition of causal models of physical processes through natural language descriptions.

Janet L. Kolodner is a professor in the College of Computing at the Georgia Institute of Technology. She received her Ph.D. in computer science from Yale University in 1980. Her research investigates issues in learning, memory, and problem



solving. As part of these investigations, she pioneered a reasoning method called case-based reasoning. Kolodner wrote the book *Retrieval and Organizational Strategies in Conceptual Memory: A Computer Model* and edited *Memory, Experience, and Reasoning. Proceedings: Case-Based Reasoning Workshop* is a collection of papers describing the state of case-based reasoning in 1988. She is currently working on a case-based reasoning textbook and has authored dozens of technical papers.



Daryl T. Lawton is an associate professor in the College of Computing at the Georgia Institute of Technology. He is also a senior research scientist in the Image Processing Branch of the Electro-Optics Laboratory at the Georgia Institute of Technology. He received a Ph.D. in 1984 in computer vision from the University of Massachusetts at Amherst, where he was also a postdoctoral fellow. From 1985 to 1989, he was with Advanced Decision Systems in Mountain View, California, where he was a principal research scientist and robot vision program manager. His current research interests involve processing dynamic image sequences, mobile robots, perceptual organization, image-understanding environments, model-based vision systems, and hypermedia for education.



Ashwin Ram received his bachelor of science in electrical engineering from the Indian Institute of Technology, New Delhi, in 1982, and his Master's in computer science from the University of Illinois at Urbana-Champaign in 1984. He received his Ph.D. degree in 1989. Ram is currently an assistant professor in the College of Computing at the Georgia Institute of Technology. He has published several papers in the areas of machine learning, natural language understanding, abduction and explanation, and cognitive science. He is on the board of editors of *Applied Intelligence* and on the technical program committee for the 1991 International Joint Conference for Artificial Intelligence.