AI Magazine Volume 8 Number 1 (1987) (© AAAI)

RESEARCH IN PROGRESS

Mallory Selfridge, Donald J. Dickerson, and Stanley F. Biggs

Cognitive Expert Systems and Machine Learning: Artificial Intelligence Research at the University of Connecticut

Abstract In order for next-generation expert systems to demonstrate the performance, robustness, flexibility, and learning ability of human experts, they will have to be based on cognitive models of expert human reasoning and learning We call such next-generation systems cognitive expert systems. Research at the Artificial Intelligence Laboratory at the University of Connecticut is directed toward understanding the principles underlying cognitive expert systems and developing computer programs embodying those principles The Causal Model ACquisition System (CMACS) learns causal models of physical mechanisms by understanding real-world natural language explanations of those mechanisms. The Going Concern Expert (GCX) uses business and environmental knowledge to assess whether a company will remain in business for at least the following year The Business Information System (BIS) acquires business and environmental knowledge from in-depth reading of real-world news stories. These systems are based on theories of expert human reasoning and learning, and thus represent steps toward next-generation cognitive expert systems.

Research at the Artificial Intelligence Laboratory of the University of Connecticut is currently focused on a number of projects addressing both fundamental and applied aspects of next-generation expert systems and machine learning. We believe that these next-generation expert systems will have to be based on cognitive models of expert human reasoning and learning in order to perform with the ability of a human expert. Consequently, we term such next-generation expert systems *cognitive expert systems.*

Cognitive expert systems should display three characteristics. First, because expert human reasoning and learning rely in part on qualitative causal models and large-scale event-based memory structures, cognitive expert systems should rely on similar knowledge. Second, because human experts are skilled at acquiring knowledge, often through natural language interaction, cognitive expert systems should learn through real-world natural language interaction. Third, cognitive expert systems should not merely learn from the input provided but should be aware of the information that has not been provided and ask the user for it. (In the extreme case, interacting with such a system consists entirely of answering its questions.) Thus, we seek to build cognitive expert systems that model expert human reasoning, learn from real-world natural language interaction, and ask questions about what is not understood.

To achieve these goals, our research falls into three major areas: (1) diagnostic expert systems that learn causal models for physical mechanisms by understanding real-world natural language explanations, (2) expert systems that judge performance of corporations and learn about them by reading real-world sources, and (3) fundamental research in computer models of cognitive development.

Mallory Selfridge is an assistant professor in the Department of Computer Science and Engineering, Donald J Dickerson is a professor in the Department of Psychology, and Stanley F. Biggs is an associate professor in the Department of Accounting at the University of Connecticut, Storrs, Connecticut 06268

CMACS: Learning Causal Models of Physical Mechanisms by Understanding Real-World Natural Language Explanations

The causal model acquisition system (CMACS) (Daniell 1985; Klimczak 1986; Selfridge, Daniell, and Simmons 1985) addresses how an expert system can learn causal models of physical mechanisms by understanding real-world natural language explanations of these mechanisms. Following research conducted by deKleer and Brown (1983; 1984), CMACS represents physical mechanisms as combinations of components whose behavior it already knows about. This component knowledge consists of semantic component frames and associated knowledge of component behavior, which are similar to deKleer and Brown's confluences.

In order to learn a causal model of a new physical mechanism from an explanation, CMACS translates each explanatory statement into a representation of its meaning. If the meaning expresses a direct causal connection, CMACS builds the appropriate piece of causal model. If the statement expresses causality between two components that cannot directly be connected to each other, CMACS infers the intermediate components. If the causality expressed by the statement is ambiguous, CMACS uses its knowledge of plausible component interconnections to infer the correct causality and build the correct piece of causal model. If the statement refers to a subcomponent of another component, CMACS merges the subcomponent information with the other component in the model. After processing the entire explanation, CMACS examines the causal model it has built and searches for components that have inputs or outputs which are not connected to other components. Such incompletely connected components represent areas of ignorance for CMACS; it generates natural language questions about these areas and builds additional pieces of the causal model from the user's answers.

After acquiring a causal model for a mechanism, CMACS can use the model to answer questions and generate explanations about the mechanism and to understand mechanism behavior and diagnose mechanism failures. CMACS answers questions and gives explanations by expressing elements of the model through a natural language generator (Cullingford et al. 1982). It understands mechanism behavior and diagnoses failures by (1) reasoning backward through the model from external behavior (for example, gauge behavior) to hypothesize possible control changes or component failures that might cause the behavior, (2) performing qualitative simulation of each hypothesized control change or mechanism failure to generate predictions about subsequent gauge behavior, and (3) comparing predicted gauge behavior with actual behavior to diagnose the control change or component failure.

CMACS currently knows 45 different components and has a vocabulary of 175 words. It successfully learns causal models of an air conditioner and a gas turbine by reading explanations taken from repair and maintenance manuals. After learning, CMACS can answer questions and generate explanations about each mechanism and can understand mechanism behavior and diagnose mechanism failures within the limits of its qualitative reasoning abilities. Our current research with CMACS involves improving these abilities by (1) developing techniques for qualitative reasoning about the behavior of compressible fluids, the relationship between passive components and active components in a mechanism, and various types of mechanical motion; (2) expanding the vocabulary, component knowledge, and learning abilities to include a much wider variety of explanations; and (3) exploring CMACS's possible role in an intelligent tutoring system.

GCX: Using Event-Based Memory Structures in an Expert System to Make Going-Concern Judgments

The going-concern expert (GCX) (Biggs and Selfridge 1985, 1986a, 1986b; Selfridge, Biggs, and Krupka 1986) is being developed to address whether a company is a going concern, that is, whether it will remain in business for at least the following year. The question is important because it is a fundamental part of an auditor's certification of a company's financial statements. If, for example, auditors determine that a company is not a going concern, then generally accepted accounting principles are not strictly applicable. Interviews with auditors have revealed that not only do they use knowledge of financial performance in making going-concern judgments, but they also have extensive knowledge of business and environmental factors which represent the underlying causes of financial performance. Thus, building GCX has involved developing both a financial-reasoning capability and a capability to reason causally from business and environmental factors. The going-concern judgment is an appropriate research topic not only because it is an important part of auditing but also because it is representative of a large set of similar problems, both within auditing and within the larger context of understanding organizational behavior.

GCX begins with financial data from a company's annual reports for a period of time and with business and environmental knowledge covering this same period. The business and environmental knowledge consists of a network of business and environmental events linked by temporal and causal relationships. In addition, this knowledge is organized and indexed through a set of higher-level knowledge structures, similar to Schank's memory organization packets (MOPS) (Schank 1982), that capture business and environmental goal structures and generalizations. Finally, causal links from business and environmental knowledge to a financial model of the company are maintained to enable GCX to understand the financial effects of business and environmental events.

GCX uses a set of financial reasoning rules to calculate a number of financial measures and evaluate these measures to make judgments about the company's financial performance. It then accesses its knowledge of the company's business and environment and uses it in a series of processing steps to refine its judgments. First, it understands the underlying reasons for the company's financial performance. It might understand, for example, that a company's recent operating loss was caused by an increase in fuel costs. Second, it uses this knowledge to judge whether the company's plans to overcome unfavorable conditions will succeed. For example, if told that the company plans a certain action, GCX examines its knowledge of the company to see if a similar action has ever been taken before and, if it has been, to assess whether it was successful. Finally, based on its evaluation of the company's financial state and its judgments of the company's plans to solve its underlying problems, GCX judges whether the company is a going concern. It can then answer questions about its judgment and the underlying problems faced by the company.

To date, GCX has been tested on data from five different companies and has extensive business and environmental knowledge about three. It has approximately 100 financial-reasoning rules, and its largest knowledge base consists of about 100 business and environmental events causally and temporally linked. GCX makes going-concern judgments about all five companies that are qualitatively similar to those of expert auditors. It uses business and environmental knowledge to reason about the possible success of the company's plans in a manner similar to expert auditors, and it generates explanations for company performance that are similar to the explanations of expert auditors.

BIS: Acquiring Event-Based Memory Structures

Because systems such as GCX require vast amounts of business and environmental knowledge, the business information system (BIS) (Moreland 1985; Moreland and Selfridge 1986), was developed to address acquiring such knowledge from in-depth reading of real-world news stories. It deals with a number of research issues, including the integration of natural language and memory structures to understand large and complex stories in detail and to make *cross-story inferences*, the "putting two and two together" types of inferences that cannot be made from information in a single story but rather must be made on the basis of information in two or more stories.

In order to acquire business and environmental knowledge through an in-depth reading of news stories, BIS reads each story one sentence at a time. It constructs a meaning representation for each sentence, unifying concepts occurring in the sentence with concepts already in the knowledge base if they are the same. BIS then examines a number of inference rules that embody knowledge of how to build the memory structure from the meanings of input sentences. These rules are used to create the appropriate links between the input meaning and existing concepts in the memory structure, to infer the existence of missing concepts, and to modify concepts on the basis of new information. Once the meaning of the input has been added to the knowledge base, other inference rules are examined. These rules are tested against the knowledge base, and conclusions are drawn on the basis of learned information. If the information triggering such inferences was originally derived from different stories, then the inference made is a cross-story inference. After a set of news stories has been processed, BIS is able to answer questions and generate summaries and paraphrases about story content and its inferences and conclusions.

. . . research on the modeling of children's cognition and learning . . . will lead to fundamental insights into the nature of expertise and learning . . .

BIS currently has a vocabulary of 300 words and has been tested on two sets of three stories about the telecommunications industry, with each set consisting of about 35 fifteen-word sentences. It successfully understands both sets of stories, performing all the required memory accesses and entity merging, and builds the appropriate memory structures, including the results of cross-story inferences. The primary results of our research are (1) the development of an architecture and a set of inference rules with which BIS learns a memory structure by reading real-world news stories in depth and (2) the development of a number of representational structures and techniques that together with those of the GCX system represent a promising approach to the representation of business and environmental knowledge. Thus, BIS represents a powerful approach to the problem of knowledge acquisition by cognitive expert systems.

Our current research with GCX and BIS involves (1) exploring the feasibility of unifying GCX and BIS into a system that acquires business and environmental knowledge, (2) expanding the business and environmental knowledge of both systems and refining it to closely match the cognitive characteristics of human experts, (3) adding the ability to generate queries about unexplained business and environmental events, and (4) extending the systems' knowledge to encompass a variety of other domains requiring cognitive expertise.

Cognitive Development

In addition to research previously described that is intended to address issues of direct and long-term relevance to various application areas, we are also pursuing research on the modeling of children's cognition and learning. We believe that such work will lead to fundamental insights into the nature of expertise and learning. Our efforts in this regard are concentrated in the area of child language learning, with secondary projects in several other areas.

Because an understanding of how children acquire language is important for theories of learning, we are carrying out long-term research on child language acquisition with the CHILD program (Selfridge 1980, 1986a, 1986b). We expect that this research will not only add to knowledge about child language acquisition per se but will also have direct implications for adult-level natural language processing by next-generation expert systems, particularly in the areas of vocabulary acquisition and overall robustness. CHILD is equipped with a natural language analyzer and generator, language learning algorithms, and an inference mechanism for using context to infer the meanings of incompletely understood utterances.

Currently, CHILD begins with the type of world knowledge possessed by a typical one-year-old and, upon receiving language input which is similar to that received by children, learns to understand and generate active and passive declarative and imperative statements in a developmental progression which is similar to that followed by children between the ages of one and five years. Specifically, it manifests language behavior consistent with six facts of language development, including the making of certain types of language errors and the subsequent recovery from the making of these errors. In addition, CHILD learns to recognize utterances as grammatically well formed. This recognition occurs not as a specific task but as a side effect of the learning that produces the developmental progression. CHILD is relevant to theoretical issues in language acquisition because it makes a number of new predictions about child language acquisition and because the underlying theory provides answers to a number of classic questions about child language acquisition. Thus, we believe CHILD represents a powerful approach to modeling child language acquisition.

In addition to the CHILD program, we are carrying out several other secondary research projects on building computer models of various aspects of child learning. The COUNT program (Selfridge and Selfridge 1986) models how children learn to count. COUNT learns to count from psychologically plausible input and makes counting errors typical of children. The IMP program (Dickerson 1986) models sensorimotor learning typical of infants and incorporates a model of the motivation behind such learning. The BALANCE project, just under way, is a computer program which learns to make judgments about the Piagetian "balance scale" and which manifests a developmental progression similar to that of children. Finally, we are using concepts developed in the domain of robot plan learning and plan recognition to model observational learning by children (Selfridge and Dickerson 1985).

Other Research

In addition to our primary research efforts, we are also carrying out or have carried out smaller-scale research projects in a number of other areas. The TARGET system (Guzzi 1985; Selfridge and Guzzi 1986) operates with a real-world dual-manipulator robot workstation and learns robot assembly plans by performing plan recognition on the demonstration of an assembly task. The NLR system (Selfridge and Vannoy 1986) operates with the same real-world dual-manipulator robot workstation and learns robot assembly plans on the basis of natural language interaction with a user. The MEMEX system is a rule-based expert system for diagnosing marine diesel engine faults that incorporates a natural language interface for communication with the user (MacDonald, Dickerson, and Selfridge 1986). The INVENTOR system (Selfridge and Cuthill 1986) addresses questions of creative invention and computer-assisted engineering. INVENTOR accepts a set of I-O specifications and then infers the causal model for a mechanism that fulfills the specifications.

Future Directions

The Artificial Intelligence Laboratory of the University of Connecticut is engaged in a program of research in cognitive expert systems and machine learning. We intend to continue this research and will concentrate on applications of cognitive expert systems and knowledge acquisition to real-world contexts as well as further our fundamental research and the transfer of its results to real-world applications.

The Artificial Intelligence Laboratory

The Artificial Intelligence Laboratory at the University of Connecticut is part of the Computer Applications Research Center, which is part of the University of Connecticut School of Engineering. It currently includes Mallory Selfridge, Donald J. Dickerson, Stanley F. Biggs, Benjamin Moreland, Anthony Guzzi, George Krupka, Barbara Cuthill, Saundra Wallfesh, and Peter Gattinella and has been in existence since 1984. The University of Connecticut artificial intelligence (AI) curriculum includes both undergraduate and graduate courses in AI programming, AI, natural language, expert systems, robotics, and machine learning.

Acknowledgments

Partial research funding has been provided by the state of Connecticut; United Technologies Research Center; the Peat Marwick Mitchell Foundation; and Ship Analytics, Inc.

References

Biggs, S F , and Selfridge, M. 1986a Developing an Expert System for the Auditor's Going-Concern Judgment Papei presented at the Fourth Annual Audit Symposium at the Univeisity of Southern Califoinia

Biggs, S. F , and Selfiidge, M. 1986b Combining Rule-Based and Network-Based Logic in Expert Systems for Accounting Paper piesented at the Ameiican Accounting Association Annual Meeting

Biggs, S F., and Selfridge, M 1985 Developing an Expert System for the Auditoi's Going-Concern Judgment Paper presented at the Joint National Meeting of The Institute for Management Science and the Operations Research Society of America.

Cullingford, R E ; Krueger, M W ; Selfiidge, M.; and Bienkowski, M 1982 Automated Explanations as a Component of a Computer-Aided Design System. In *IEEE Transactions of Systems, Man, and Cybernetics* 12(2): 168–181

Daniell, J. D 1985 Leaining Causal Models by Undeistanding Real-World Natuial Language Explanations M S thesis, Technical Report, #TR86-21, Dept of Computer Science and Engineering, Univ of Connecticut

deKleer, J , and Brown, J S , 1984 A Qualitative Physics Based on Confluences *Artificial Intelligence* 24: 7–84 deKleer, J , and Brown, J S , 1983 The Origin, Form, and Logic of Qualitative Physical Laws In Proceedings of the Eighth International Joint Conference on Artificial Intelligence, 1158–1169. Menlo Park, Calif : American Association of Artificial Intelligence

Dickerson, D. J 1986 IMP: A Computer Model of Sensorimotor Motivation and Learning, Technical Report, #TR86-20, Dept of Computer Science and Engineering, Univ of Connecticut

Guzzi, A 1985 Learning Robot Assembly Plans by Observation M S thesis, Technical Report, #TR86-22, Dept. of Computer Science and Engineering, Univ of Connecticut.

Klimczak, B 1986. Learning Causal Models of Physical Mechanisms from Real-World Explanations Using a Global Dictionary. M S thesis, Technical Report, #TR86-23, Dept of Computer Science and Engineering, Univ of Connecticut

MacDonald, C.; Dickerson, D J ; and Selfridge, M. 1986 Natural Language Processing within a Rule-Based Expert System for Diagnosing Marine Diesel Engine Faults, Technical Report, #TR86-24, Dept. of Computer Science and Engineering, Univ of Connecticut

Moreland, B , and Selfridge, M 1986 Acquiring Business and Environmental Knowledge from Real-World Natural Language Text, Technical Report, #TR86-15, Dept of Computer Science and Engineering, Univ of Connecticut

Moreland, B 1985. Cross-Story Inferences Made during Competitive Analysis: A Computer Model M S thesis, Technical Report, #TR86-25, Dept. of Computer Science and Engineering, Univ of Connecticut

Schank, R C. 1982 *Dynamic Memory: A Theory of Learning in Computers and People*. Cambridge, Mass.: Cambridge University Press.

Selfridge, M 1986a A Computer Model of Child Language Acquisition *Artificial Intelligence* 29: 171–216

Selfridge, M. 1986b How Children Learn to Recognize Grammaticality: A Computer Model, Technical Report, #TR86-16, Dept of Computer Science and Engineering, Univ. of Connecticut.

Selfridge, M ; Biggs, S F ; and Krupka, G 1986 GCX: A Cognitive Expert System for Making Going-Concern Judgments, Technical Report, #TR86-20, Dept of Computer Science and Engineering, Univ. of Connecticut

Selfridge, M., and Cuthill, B. 1986 Automatic Invention: Inferring the Causal Model of a Physical Mechanism from a Description of Desired Behavior, Technical Report, #TR86-17, Dept. of Computer Science and Engineering, Univ of Connecticut.

Selfridge, M , and Guzzi, A. 1986. Learning Robot Assembly Plans by Observation, Technical Report, #TR86-18, Dept of Computer Science and Engineering, Univ of Connecticut

Selfridge, M , and Selfridge, O G. 1986 How Children Learn to Count: A Computer Model, Technical Report, #TR86-19, Dept of Computer Science and Engineering, Univ of Connecticut

Selfridge, M., and Vannoy, W 1986 A Natural Language Interface to a Robot Assembly System *IEEE Journal of Robotics* RA-2(3).

Selfridge, M ; Daniell, J. D ; and Simmons, D. 1985 Learning Causal Models by Understanding Real-World Natural Language Explanations In Proceedings of the Second International Conference on Artificial Intelligence Applications, 378–383. Washington D C : IEEE Computer Society

Selfridge, M , and Dickerson, D 1985. Observational Learning by Computer In Proceedings of the Second International Conference on Artificial Intelligence and Education, 95–97

Selfridge, M 1980 A Process Model of Language Acquisition Ph D diss , Research Report, #172, Dept of Computer Science, Yale Univ