

THE EMERGENCE OF ARTIFICIAL INTELLIGENCE TECHNOLOGY

Ronald J. Brachman and Fran H. Henig

AT&T TECHNICAL JOURNAL

Ronald J. Brachman is head of the Artificial Intelligence Principles Research Department at AT&T Bell Laboratories in Murray Hill, New Jersey. Before joining AT&T in 1985, he participated in designing and implementing two well-known knowledge-representation systems, KL-ONE and KRYPTON. Mr. Brachman received a B.S. in electrical engineering from Princeton University and Master's and Ph.D. degrees in applied mathematics from Harvard University. **Fran H. Henig** is head of the Artificial Intelligence Systems Department of Bell Laboratories in Liberty Corner, New Jersey, responsible for applying AI technology to the development of Operations Systems. Ms. Henig joined AT&T in 1964 with an A.B. in mathematics from Wheaton College, Norton, Massachusetts, and has an M.S. in mathematics and (continued on page 6)

This issue of the *AT&T Technical Journal* provides a sample of artificial intelligence projects under way at AT&T. These efforts extend to research and development, and cover topics from planning and language understanding to VLSI design and sonar interpretation. The breadth of the artificial intelligence field and several important applications are highlighted.

The Scope of Artificial Intelligence

Almost everyone has an opinion about *artificial intelligence* (AI). Either it's the wave of the future, or it's just a public relations canard. It's a new programming technology or it's a new mindset—a way of looking at problems that no other discipline has. Or, perhaps the name is an oxymoron and the whole idea of AI is absurd.

The popular press has recently reported that "AI is here."¹ But, to be honest, it has by no means totally arrived. On the other hand, such statements are not pure hyperbole either (although AI can boast its share of that).

Despite the disparity of attitudes toward AI, there are indeed new and substantive results emerging—results that are important scientifically and worthy of attention from corporations such as AT&T.

In fact, AI is already a part of some AT&T products, and promises some important developments in the next few decades.

Our goal in this issue is to give you a snapshot of the AI work at AT&T. As you thumb through these articles, you will begin to realize the scope of work in the field. Perhaps this is one source of misunderstanding about AI. AI is extremely broad; it seeks to understand and, perhaps ultimately, emulate the entire scope of activity (typically human) that we can comfortably label "intelligent." Given this, it is easy to see that AI is not a simple commodity to be folded generously into a development project in an attempt to turn a gizmo into an intelligent gizmo.

It is the study of a large variety of behaviors that we consider intelligent, whether or not they share any underlying foundation. Currently, however, AI is not the study of all of these at once. Rather than pursuing the commonly attributed (but unbelievably ambitious) goal of

building a complete, sentient, conscious, autonomous robot, researchers in AI are looking at the more prosaic, small components out of which intelligent behavior seems to arise.

We would perhaps be better off using the term “computational reasoning” to represent what we do. Equipped with insights from computer science, we are attempting to look at problems of reasoning and rationality not addressed previously by disciplines other than psychology and philosophy.

In AI we use the computational metaphor to help us find answers to questions such as:

- How can an intelligent system anticipate the future and prepare for it, without wasting effort on irrelevant or extremely improbable events?
- How can we make systems that function with only an impoverished, partial knowledge base in the way that humans are able to do?
- How can we communicate successfully with computers in (natural) languages fraught with ambiguity and vagueness?
- How can multiple entities plan and cooperate to achieve goals that no individual could ever hope to?

While we can look to human behavior for insight into these issues, our main goals are to understand the phenomena of intelligent activity and find ways to exploit that understanding in artifacts.

We are far from being able to answer these questions, or even to integrate what little we know about intelligent behavior. Yet, there has been some interesting progress in the last thirty years, and the efforts at AT&T are growing in prominence and effectiveness. This issue of the *AT&T Technical Journal* should give the interested reader a chance to see what we can do now, as well as some insight into some of the more interesting open issues of the day. There is much in the field to steer clear of, but we think that you will see from this collection of papers that AT&T is pursuing this somewhat enigmatic discipline in a sober and serious way and with tangible and previously unexpected results.

In This Issue

Because artificial intelligence studies such a broad range of phenomena, it is impossible to present the state of the art in its entirety. (The recently published *Encyclopedia of Artificial Intelligence*² has over 1100 pages.)

Research is progressing in most areas, with important applications forthcoming from a few. The important areas for research include:

- Knowledge representation
- Commonsense reasoning
- Expert systems
- Machine learning
- Natural language understanding and generation
- Computer vision (and other forms of perception)
- Planning, design, diagnosis, and other forms of problem-solving.

AI is also concerned with hardware and software systems to support these research areas (e.g., a recent, serious foray into highly parallel architectures, including those of the “connectionist” or “neural network” type). And it also intersects broadly with work on robotics and speech recognition, both areas of prominence within AT&T. AT&T has become a leader in some of the core areas of AI research. Most of our strengths are represented in this issue of the *AT&T Technical Journal*.

In the first article, Brachman³ traces the central area of knowledge representation and reasoning. This is a tutorial paper that explains the role of mechanized reasoning in AI and surveys the most common techniques used to represent knowledge in current AI systems. As Brachman points out, knowledge representation is fundamental to virtually all work in AI. This article discusses why and how.

Kautz and Pednault⁴ give a clear explanation of work involving the synthesis and analysis of plans. Almost since its inception, AI has been studying automatic planning, first in the context of robot tasks and, more recently, in the context of more sophisticated enterprises such as autonomous deep-space missions and the coordination of factory floor activities. The other side of planning has also become important. This is the area of plan recognition (the

determination of what set of goals an agent that you are observing is attempting to satisfy). This will be a crucial ingredient of complex natural-language conversations and interfaces to complex systems, such as computer operating systems and network operations systems.

Hirschberg, Ballard, and Hindle⁵ describe their work in natural language processing. A brief overview of the concerns of this prominent area precedes discussions of projects in syntactic processing, use of intonation information in speech generation, and a transportable natural-language "front end" that improves interaction with formal systems such as databases. Natural language has been a primary focus of artificial intelligence for many years, and represents one of AT&T's greatest strengths.

Howard, Jackel, and Graf,⁶ assisted by Jones,⁷ convey some of the burgeoning excitement about neuro-morphic architectures, or "artificial neural networks" in their companion articles. Rather than working on higher-level, more cognitive aspects of human reasoning and problem solving, many researchers are looking at more direct emulation of some of the apparent information-processing structures of the brain. These articles describe some of the leading research on the fabrication of chips with many neuromorphic structures and some of the processing of which these chips seem capable.

Work on expert systems has recently dominated AI-related developments. These are systems that attempt to capture human expertise in very narrow technical areas. AT&T has been a pioneer in expert systems; the ACE (Automated Cable Expertise) system⁸ was one of the first expert systems to be sold commercially and today is one of the most widely deployed. Several of the papers in this issue describe recent projects that are attempting to exploit AI technology in other important applications.

Vesonder⁹ discusses some of the tools currently available within AT&T for building expert systems. His discussion centers on the rule-based programming environment. C5, a locally-developed, C-based variant of OPS5¹⁰ is highlighted, and other portable, UNIX[®] system-based tools are described. More generally, Vesonder dis-

cusses the issue of integrating expert systems technology with more standard programming technologies.

Following that, three important expert systems projects are presented. For now, these are the types of AI systems that will have the most direct impact on our lives, both internally as aids to productivity, and externally as products. First, Kowalski¹¹ discusses his extended work on developing a system that can emulate an expert in designing microcomputers. This system, the Design Automation Assistant (DAA), takes an algorithmic description of a VLSI (very-large-scale integration) system and produces a list of technology-dependent registers, operators, data paths, and control signals. The DAA uses knowledge gleaned from many experts during many hours of interviews. Kowalski addresses the design process in general, gives details of the construction of the DAA, and briefly discusses the system's performance, especially in the context of its successful design of an IBM System/370.

Callahan¹² then presents TOPAS-ES (Testing Operations Provisioning Administration System-Expert System), a system used to improve network maintenance. TOPAS-ES is intended for use in AT&T Trunk Operations Centers. The system is among the first of a generation of "real-time, distributed, multi-tasking expert systems." Callahan outlines the complexities of the network maintenance task, and examines how part of the job can be done by two expert systems working together cooperatively (the trouble analyzer and the trouble sectionalizer).

Finally, Kuzmak, Schoenfelt, Gorin, Brittingham, and Milich¹³ give an impressive account of a knowledge-based approach to signal interpretation. In this paper, we see an example of how AI technology can be successfully integrated and used with more standard approaches to problems. Here, hybrid representation technology and parallel processing are used to address a complex interpretation problem (grouping detections believed to originate from the same platform) and to classify the contacts. This appears to represent the most sophisticated extant approach to this problem, and shows again how mainstream and AI technologies can mingle productively.

References

1. "AI is Here," *Business Week*, No. 2850, July 9, 1984, pp. 54-62.
2. *Encyclopedia of Artificial Intelligence*, S. C. Shapiro, ed., John Wiley & Sons, New York, 1987.
3. R. J. Brachman, "The Basics of Knowledge Representation and Reasoning," *AT&T Technical Journal*, Vol. 67, No. 1, January/February 1988, pp. 7-24.
4. H. A. Kautz and E. P. D. Pednault, "Planning and Plan Recognition," *AT&T Technical Journal*, Vol. 67, No. 1, January/February 1988, pp. 25-40.
5. J. Hirschberg, B. W. Ballard, and D. Hindle, "Natural Language Processing," *AT&T Technical Journal*, Vol. 67, No. 1, January/February 1988, pp. 41-57.
6. R. E. Howard, L. D. Jackel, and H. P. Graf, "Electronic Neural Networks," *AT&T Technical Journal*, Vol. 67, No. 1, January/February 1988, pp. 58-64.
7. M. A. Jones, "Programming Connectionist Architectures," *AT&T Technical Journal*, Vol. 67, No. 1, January/February 1988, pp. 65-68.
8. F. D. Miller, J. R. Rowland, and E. M. Siegfried, "ACE: An Expert System for Preventive Maintenance Operations," *AT&T Bell Laboratories Record*, January 1986, pp. 20-25.
9. G. T. Vesonder, "Rule-Based Programming in the UNIX® System," *AT&T Technical Journal*, Vol. 67, No. 1, January/February 1988, pp. 69-80.
10. L. Brownston et al., *Programming Expert Systems in OPS5: An Introduction to Rule-Based Programming*, Addison-Wesley, Reading, Massachusetts, 1986.
11. T. J. Kowalski, "The VLSI Design Automation Assistant: A Synthesis Expert," *AT&T Technical Journal*, Vol. 67, No. 1, January/February 1988, pp. 81-92.
12. P. H. Callahan, "Expert Systems for AT&T Switched Network Maintenance," *AT&T Technical Journal*, Vol. 67, No. 1, January/February 1988, pp. 93-103.
13. S. D. Kuzmak et al., "Knowledge-based Signal Interpretation," *AT&T Technical Journal*, Vol. 67, No. 1, January/February 1988, pp. 104-120.

Biographies (continued)
computer science from the Stevens Institute of Technology.

(Manuscript received November 19, 1987)

JANUARY/FEBRUARY 1988 • VOLUME 67 • ISSUE 1

AT&T TECHNICAL JOURNAL