

# R1 And Beyond:

## AI Technology Transfer At DEC

**Stephen Polit**

*Digital Equipment Corporation, 77 Reed Rd., Hudson, MA 01749*

---

### Abstract

This article describes one person's experience in coming from an academic environment to work at Digital Equipment Corporation. The author feels his own experience has paralleled the transfer of AI technology from academia to industry, where AI researchers must live up to very different expectations, but also enjoy very different rewards. This article covers the historical background of DEC's involvement with AI, the development of R1—known internally and henceforth in this article as XCON—and DEC's experiences with it and its consequences. Finally, the article offers advice for other corporations planning to develop their own capabilities in AI.

---

### DEC and AI—A Historical Perspective

DEC is unlike most other corporations, even most other computer corporations, in that it has a fairly long and involved historical relationship with AI. I won't give all the details of that history. Rather, I will explain how that history has affected perceptions and attitudes at DEC and other corporations, and how those attitudes have influenced DEC's decisions in the area of artificial intelligence. I've divided this history into two distinct parts. The first part describes DEC's years as the main hardware vendor for the AI community, and the second part describes DEC's more recent role as a designer and user of AI software.

DEC was the main hardware vendor for the AI community during the 1970s, a period characterized by the birth and development of expert systems in academic research laboratories. Since DEC had a very close relationship with the academic computing laboratories, much of this work was done on the DECSYSTEM-10 and its successor, the DECSYSTEM-20. The two major dialects of LISP, MACLISP and INTERLISP, were developed on these machines; so were most of AI's early successes and disappointments. This facilitated an exchange of engineering talent between university AI labs and DEC's engineering labs. Thus, even before the commercialization of AI and DEC's entry into the AI market, DEC engineers were familiar with AI software projects and techniques. As AI became increasingly commercial, DEC's history of involvement

with AI led people both inside and outside DEC to believe that DEC was a natural supplier of machines and tools for the development of AI systems. Consequently, DEC was open to entering the AI software market very quickly.

I'll discuss XCON in a bit more detail later on, but right now I'll describe how its success affected attitudes towards AI within and outside DEC. XCON was originally developed at Carnegie-Mellon University in 1979 by John McDermott and several associates. It was one of many DEC-university collaborations; DEC profited from almost all of these, as did the universities. One benefit of this project was that a closer tie was created between DEC and a major AI lab. Although the system received DEC's moral and financial support during its development, it did not immediately represent a large commitment of DEC to artificial intelligence. The successful completion of R1 and its incorporation as XCON, however, changed attitudes towards AI within DEC. The demonstrated practicality of AI made it a more legitimate subject of industrial interest. Also, somewhat ironically, it suddenly gave DEC the image of being a leading practitioner of AI, a prominent example of an industrial firm successfully using an AI tool. Finally, and importantly, considerable experience was gained by bringing XCON to use at DEC.

### Building Expert Systems

I'll now give a brief review of the steps involved in building expert systems as they are described by many researchers. The five steps involved in building an expert system are:

- Step 1: problem recognition,
- Step 2: task definition,
- Step 3: initial design,
- Step 4: knowledge acquisition, and
- Step 5: system maintenance.

First, during step 1, someone must recognize that there is a problem to be solved and determine whether AI is an appropriate way of solving it. Frequently, the problem is perceived as a bottleneck in a larger process; sometimes it is a scarcity of trained personnel. Second, during step 2, researchers must define the functions the AI system will perform. These two steps may be the most difficult ones, because researchers must have a comprehensive understanding

---

This article is an edited version of Dr. Polit's presentation at the Technology Transfer Symposium held at the AAAI-83 conference

of all available AI tools and techniques as well as of the problem itself.

During step 3, initial design, choices of particular AI techniques, an inference mechanism, and knowledge representations are made. Then, during step 4, the knowledge acquisition phase, the expertise of people performing the task is entered into the knowledge base of the system. In fact, however, steps 2, 3, and 4 do not take place sequentially; there is an interaction between them. During the knowledge acquisition phase, a greater understanding of the problem is obtained. The design of the system may be changed because unexpected difficulties are recognized, or the task may be redefined because perception of the problem has changed.

The final step, system maintenance, is perhaps the most important one. It probably represents the largest cost of the system—keeping it running smoothly, extending it, refining it, and adding new knowledge. Experts learn things, and the people maintaining an expert system should be able to make that system learn some of these things as well. A knowledge system that is not updated will quickly become obsolete.

### The XCON Experience

I will now relate the expert-system-building steps to the actual experiences DEC had with building XCON. XCON is an expert system that expands and validates orders for DEC VAX-11 computer systems and then configures the placement of components within the equipment cabinets.<sup>1</sup>

Customers purchasing DEC computers may select from a wide variety of components; the problem is that some orders then contain combinations of pieces that will not operate together, and some orders lack essential components. Before XCON was designed, technical editors examined the orders, identified nonfunctional combinations of components, and provided instructions for assembling the system. Since this process caused a bottleneck at DEC, and many attempts to automate it had failed, DEC was willing to try to solve this problem with a technology that was relatively unexplored in an industrial setting.

John McDermott, then a research associate at Carnegie-Mellon University, felt that DEC's configuration problem could be solved using a rule based expert system. McDermott and his associates performed the task definition, system design, and knowledge acquisition steps of expert system development at Carnegie-Mellon, although knowledge acquisition required much contact with DEC technical editors. The initial XCON was delivered to DEC in 1980. It was a large system, containing about 750 rules, and even though it could configure many of the orders submitted to it, it required extension and refinement by DEC before it could become the dependable system it is today.

DEC learned that it could not simply use the system as designed in the university and keep it running. The system's knowledge had to be corrected and extended by people

trained in knowledge acquisition and expert systems programming. A group was formed at DEC to assume responsibility for maintaining XCON. Unfortunately, its members were at a severe disadvantage. Most of them were DEC engineers who were not expert at AI. Even though some DEC engineers were familiar with AI, the maintenance group was not chosen for its AI expertise. Also, the system was written in OPS-4, a production system language developed at Carnegie-Mellon, with which no one in the group was familiar. After about a year and much aid from Carnegie-Mellon, the group was able to take over maintenance of XCON (which had grown to about 1,000 rules).

This was a successful but difficult experience. DEC learned the hard way that maintaining an expert system requires specially trained personnel. In doing so, it trained some of those personnel. Today, XCON has grown to 2,500 rules and is far superior to the system that was delivered in 1980. It configures orders for all VAX-11 computer systems manufactured in the United States, as well as some orders for other machines. It is continually modified and updated to accommodate new options. XCON has become an integral part of the manufacturing process, and this integration was achieved by a group of people who had no previous experience with developing expert systems. The only problem was the time required for them to learn how.

Many people feel that the primary significance of the XCON experience is that AI techniques were used to solve real-world problems in an industrial setting. However, this is less significant than the fact that an industrial, rather than an academic, user was able to successfully help implement and maintain an AI system. Finally, DEC learned that knowledge-based systems must be maintained by trained personnel.

### After XCON

Around 1982, after the successful development of several academic expert systems and the commercial success of XCON, AI was widely publicized in the popular press. The Japanese made AI the focus of their fifth-generation computing effort, and a number of American corporations became interested in developing expert systems. DEC was no exception; it entered a period of widespread expert systems activity.

One of the many systems developed at DEC during this period was XSEL, an expert system related to XCON. "SEL" stands for "seller's assistant"; XSEL aids sales representatives in tailoring a customer's order to that customer's particular needs. It then delivers the order, along with the customer-specific constraints it has gathered, to XCON for configuration.

XSEL's history parallels that of XCON. Steps 1 to 3 of system development were done at Carnegie-Mellon, and the system was then delivered to DEC. Again, employees had to be trained in expert system maintenance, because the XCON maintainers did not have enough time to maintain a second system. This experience made DEC realize that it needed

<sup>1</sup>See John McDermott, "R1: The Formative Years," *AI Magazine* Summer 1981, 21-29 and Judith Bachant and John McDermott, "R1 Revisited: Four Years in the Trenches," *AI Magazine* Fall 1984, 21-34

to establish a formal training program (which is described below).

Other projects were XSITE, a system that helps prepare a customer's site for the installation of a purchased processor, and efforts in VLSI design, hardware fault diagnosis, educational services, and office automation. This research was performed by many different groups working independently, but there were some fairly coordinated efforts as well. One group was formed in an explicit effort to decrease DEC's dependence on university researchers. This group, known as the Knowledge Engineering Advanced Development Group, was an offshoot of the XCON group; it consisted of four AI professionals. The group's function was to contract with other groups within DEC to carry out steps 2 through 4 of those groups' expert system projects. The Knowledge Engineering Advanced Development Group was given initial two-year funding with the intent that it become financially self-supporting after that time.

The Knowledge Engineering Advanced Development Group created two successful systems, one to aid in VLSI design, the other in hardware fault diagnosis. More importantly, this was a formal attempt by DEC to increase the extent of technology transfer. While techniques were developed during the XCON and XSEL projects to facilitate the minimal transfer required for system maintenance and support, the Knowledge Engineering Advanced Development Group attempted to provide the expertise for the initial stages of expert system development. This required an explicit corporate commitment to developing expert systems technology.

Even though XCON was a success, trained personnel were available to maintain expert systems, and new systems were being successfully initiated and completed in-house, the lack of a focus of activity became a critical problem at DEC. First, few highly trained AI researchers were available, and since several groups within DEC were competing for them, no one group could secure a large number of qualified personnel. Second, there was no large group of AI researchers working in one location. There was no way to coordinate the activities of the various research groups so that, for instance, they could agree on a common set of tools. Furthermore, most of the AI professionals were accustomed to working in an environment providing sophisticated and powerful computing equipment and other AI researchers with whom they could share ideas. Without a critical number of people in one location, that environment was difficult for DEC to provide.

Another problem was resource allocation. There was no central authority to decide which projects should be done and to allocate funding. As I mentioned earlier, the method chosen for the Knowledge Engineering Advanced Development Group was to obtain funding from other groups within DEC, which meant that its choice of projects depended on the budgets of those groups. This arrangement favored short-term projects with quick results, instead of longer-term development efforts that could build a foundation for a wider class of applications. It was felt that an effective AI research group should have the freedom to choose projects without undue

reliance on the budgets of other groups. Eventually a number of people, of whom I was one, began an informal process of reorganizing AI activity at DEC to achieve the goals we felt were needed for the successful transfer of AI technology.

As a result of this process, an organization was established to make informed decisions about AI research at DEC and to provide a centralized working environment. This organization, known as the AI Technology Center, includes three technical groups. Their activities are not restricted, but they can be categorized as the engineering applications and advanced development group, the manufacturing applications group, and the customer service applications group. In addition, there is a marketing group that monitors the current AI market and the public's perception of AI. There is also a formal training program to ensure a supply of qualified AI engineers to maintain the systems DEC produces. Finally, there is a coordination committee to make major decisions and to direct the course of AI activity at DEC.

DEC's artificial intelligence training program was established to avoid the maintenance problems DEC encountered when fielding XCON and XSEL. Students are first given a fourteen-week introductory course in AI, which is taught by both DEC and university AI researchers. They then enter a nine-month apprenticeship in building expert systems at either a university or a DEC AI group. The aim of this program is not to produce AI researchers, but rather to provide a pool of people familiar with AI who can support the systems DEC develops.

I am a member of the engineering applications and advanced development group, which is called the AI Technology Group. It is currently staffed with six AI professionals having graduate degrees. We work mostly on engineering applications, but we also do advanced development work. In addition to building expert systems, we are often called upon to evaluate outside tools or DEC's corporate strategies.

## Conclusions

My advice to corporations planning to transfer AI technology is:

- Carefully calculate the level of commitment you are prepared to make, and know what you expect in return.
- Understand the risks you are taking. Get informed advice.
- Know who your AI experts are, whom you can trust technically— and consult them

As AI's potential has been understood and the technology practiced, expectations about its capabilities have grown. People's feelings about AI's influence on the world and their own particular organizations have grown as well. AI is actually changing the way DEC and many other organizations do business. My only caution is to avoid inflated expectations about the capabilities of artificial intelligence and be willing to reevaluate your goals. Even revolutionary ideas and technologies do not solve every problem.