

Is CAD/CAM Ready For AI?

Salvatore J. Stolfo*
Department of Computer Science
Columbia University

1. Introduction

The field of Computer Aided Design and Manufacture (CAD/CAM) has had major successes in the last decade. With the introduction of CAD/CAM design systems in the marketplace, many people feel that CAD/CAM is no longer the promised technology of the future, but the emerging technology of the present. With the onset of strong international competition, increasing economic pressure has demanded that our nation's productivity and manufacturing quality improve dramatically. It is not surprising, therefore, that CAD/CAM has assumed a vital position in our national economic plans.

In a similar manner, the field of artificial intelligence (AI) has had equally impressive successes and advances in its technological base. The promise of AI is no longer for the future, but rather can be seen to begin making direct impact in the commercial marketplace. In this paper, we explore the possibility of merging the two technologies synergistically to improve upon the capabilities of CAD/CAM systems as well as to provide AI researchers with a very rich and critically important application area.

It would be rather easy to speculate in a very imaginative way on how a fully automated manufacturing system might look in the future, driven primarily by autonomous artificially intelligent computers. Rather, it would be more important that we concentrate on shorter term goals. Towards that end, we shall limit our discussion to the present state and capabilities of AI knowledge-based expert systems and how these rather immature but impressive systems may guide us directly in improving present-day CAD/CAM systems.

For pedagogical reasons, we will begin with a brief tutorial on expert systems focusing on one system soon to be available in the commercial marketplace. We will then present a cursory view of other application areas that have been successfully attacked by this technology. We conclude with several observations which link

*Partial support was provided by the Office of Naval Research through contract number N0014-82-K-0256


existing expert system applications with various CAD/CAM problems. Our prospectus is bright: CAD/CAM is ready for AI!

2. Current Expert Systems Technology

Knowledge-based expert systems are AI problem-solving programs designed to operate in narrow "real-world" domains, performing tasks with the same competence as a skilled human expert. Elucidation of unknown chemical compounds [1], medical diagnosis [10] and mineral exploration [3] are three of the best known examples. These systems contrast greatly with the earlier general-purpose AI problem-solvers which were typically implemented without a specific application in mind. One of the key differences is the large amounts of problem-specific knowledge encoded within present-day systems. Thus, the field of AI has redirected much of its effort away from small, general-purpose systems to large, special-purpose programs (see figure 2-1)

Expert systems have been constructed, typically, from two loosely coupled modules, collectively forming the *problem-solving engine* (see figure 2-2). The *knowledge base* contains a large collection of facts, definitions, and heuristic 'rules-of-thumb' embodying all of the relevant domain-specific information. This domain knowledge *acquired directly from a human expert* permits the program to behave as a specialized, intelligent problem-solver


Figure 2-1: From small general systems to large specialized systems



Much of the research in AI has concentrated on effective methods for representing and operationalizing human experiential domain knowledge within a computer program. The representations that have been proposed have taken a variety of forms including purely declarative-based logical formalisms, highly-stylized" rules or productions, and structured generalization hierarchies commonly referred to as semantic nets and frames. Many knowledge bases have been implemented in rule form, to be detailed shortly

The second component of the problem-solving engine of an expert system is the *inference engine* which controls the deductive process it implements the most appropriate strategy, or *reasoning* process for the problem at hand. The earliest AI problem-solvers were implemented with an iterative branching technique searching a large combinatorial space of problem states. Heuristic knowledge, applied within a *static control strategy*, was introduced to limit the search process while attempting to guarantee the successful formation of solutions. In contrast, state-of-the-art expert systems separate the control strategy from an inflexible program, and deposit it in the knowledge base along with the rest of the domain-specific knowledge

Figure 2-3: Organization of a Problem-Solving Engine



Thus, the problem-solving strategy becomes domain-dependent, and is subject to the same methods of acquisition and deductive manipulation as facts and assertions.

As is the case with nearly all of the reported experimental systems, a knowledge base and an inference engine are clearly identifiable. In order to build these systems, though, methods are needed to effectively transfer expertise from a human to a computer. Thus, the *knowledge acquisition* problem, and machine learning in general, has attracted considerable attention from AI researchers. The acquisition process, however, must be a cooperative effort between human and computer. It is difficult to imagine a human interacting with a reasoning machine unless effective explanation of the automated problem-solving behavior is provided. Thus, the *explanation problem* has also assumed an important position in AI research. However, in many systems the acquisition of the knowledge base, and the explanation of the problem-solving activities of the system are handled by ad hoc methods.


A notable exception is the *TEIRESIAS/MYCIN* system [2], which attempted to fully automate the acquisition and explanation of knowledge provided by a human expert unsophisticated in computer technology. *MYCIN* was implemented as a rule-based consultant program performing diagnosis of bacterial infectious diseases. *TEIRESIAS* is the "front-end" system providing explanation and acquisition facilities to automate the transfer of expertise from a medical doctor to the *MYCIN* problem-solving engine. Indeed, the system architecture of *TEIRESIAS/MYCIN* depicted in Figure 2-3, has served as a model for organizing current and future expert systems.

In the present paper, we will not discuss these two areas of AI research. Rather, we will focus briefly on the problem-solving aspects of one example expert system.

3. An Example: ACE

ACE, an expert system designed for Automated Cable Expertise, is an analysis program, soon to be available commercially, perusing large volumes of data tracking failure reports in a telephone network. In the following, we describe the general function and structure of the *ACE* problem-solving engine. The reader is encouraged to see [12] for a detailed description of the system.

Figure 2-3: The structure of *TEIRESIAS/MYCIN*



3.1. The Problem

In normal operation, the telephone network supports a telephone line, called a *cable pair*, from a residential or business site. A collection of pairs are bundled together to form the cables that hang from telephone poles, or reside underground. A collection of cables form a *wirecenter*. These three levels form the bulk of the local telephone network and the cable maintenance force concentrate their efforts at all three levels.

A variety of electrical faults and environmental conditions can cause failure of one or more cables or individual pairs. An important and expensive operation performed by the telephone operating companies is general maintenance and rehabilitation of these lines.

Customer generated maintenance reports provide important information for identifying "trouble spots" within the local network. In a high-density geographic area, the logging and tracking of failure reports has become an important and expensive data processing operation. Thus, many telephone companies use a conventional database system, called *CRAS*, to monitor the maintenance of the local network on a daily basis. Highly trained analysts routinely peruse impressive volumes of *CRAS* data and attempt to identify trouble spots to prevent further disruption of service to customers. However, the limited number of specialists available, and the size of the database inhibits the timely analysis and reporting of persistent problem areas which require rehabilitation. The approach of installing *ACE*, assisting management decision making, was proposed and implemented as a solution to the long-term problem of timely and accurate selection of areas for rehabilitation.

3.2. The ACE Problem-Solving Engine

Within *ACE*, the corpus of knowledge about wirecenters, *CRAS* data and commands, and analysis strategies is embodied by a *Production System* program. As has been reported by several researchers, production system representation schemes appear well suited to the organization and implementation of knowledge-based software. Rule-based systems provide a convenient means for human experts to explicate their knowledge, and are easily implemented and readily modified and

extended. Thus, it is the ease with which rules can be acquired and explained that makes production systems so attractive

3.2.1. Production Systems

In general, a *Production System* [4, 8, 9] is defined by a set of rules, or *productions*, which form the *Production Memory* (PM), together with a database of assertions, called the *Working Memory* (WM). Each production consists of a conjunction of *pattern elements*, called the *left-hand side* (LHS) of the rule, along with a set of actions called the *right-hand side* (RHS). The RHS specifies information that is to be added to (asserted) or removed from WM when the LHS successfully matches against the contents of WM.

An English language equivalent of an *ACE* production rule is presented in figure 3-1.

Figure 3-1: An Example *ACE* production.

```
IF a range of pairs in a cable have generated
   a large number of customer reports
ANDIF
   a majority of the work on those pairs was
   done in the terminal block

THEN
   look for a common address for those repairs.
```

In operation, the production system repeatedly executes the following cycle of operations:

- 1 *Match*: For each rule, determine whether the LHS matches the current environment of WM. All matching instances of the rules are collected in the *conflict set of rules*.
- 2 *Select*: Choose exactly one of the matching rules according to some predefined criterion.
- 3 *Act*: Add to or delete from WM all assertions specified in the RHS of the selected rule.

During the selection phase of production system execution a typical interpreter provides *conflict resolution strategies* based on the *recency* of matched data in WM, as well as syntactic discrimination. Rules matching data elements that were more recently inserted in WM are preferred, with ties decided in favor of rules that are more specific (i.e. have more constants) than others.

3.2.3. The ACE Knowledge Base

Although no structure is provided (or imposed) on PM by the general production system paradigm, the set of approximately 200 rules within ACE's knowledge base can be loosely organized into subsets of related rules which collectively perform the analysis.

A set of productions performs short term analysis by examining the flow of trouble reports on a daily basis. If troubles are reported for a cable that has no previous history of troubles then information is retained that indicates that this cable may soon require attention.

When new failures are reported for a cable with a history of persistent problems, ACE requests further detailed reports from CRAS, along with a list of standardized procedures used to repair the type of troubles reported. This information is used to deduce

- 1 whether the repair task done on that cable suggests that preventive maintenance may reduce future troubles.
- 2 if preventive maintenance is required then what type is likely to be appropriate.
- 3 and, if possible, where the rehabilitation should be done

Thus, ACE not only identifies trouble spots, but also suggests how to repair them.

Another portion of ACE PM contains a set of productions which know how to communicate with CRAS. Based on requests for more data generated by other analyses, these productions assemble the appropriate CRAS commands and parameters and then monitor the resulting data stream retrieved from CRAS.

Finally, a set of rules assemble the appropriate messages about the day's events recognized by the system, and call on the electronic mail facilities to deliver them to the appropriate users. ACE knows the target of each message based on the relative importance of the message to the user.

4. Other Expert System Applications

ACE is one example of the successful application of expert systems technology to a "real-world" problem domain. Many other examples exist in a wide range of application domains including science, mathematics, engineering, and medicine. We shall briefly describe several of these which may have bearing on various CAD/CAM problems.

DENDRAL [1] is an expert in one aspect of chemistry. Given an NMR mass spectrogram of some unknown chemical compound, DENDRAL successfully identifies a relatively small set of candidate molecular structures consistent with the spectral data. The main sources of knowledge for DENDRAL were acquired directly from chemists, one of whom won a Nobel prize for research in chemistry used in the DENDRAL system.

Mathematics has been tackled by MACSYMA [6], an expert system assisting mathematicians in solving mathematical problems. Using symbolic input, rather than numeric, as well as a very large body of algebraic manipulation primitives, MACSYMA has gained wide use as well as acclaim for its adeptness at simplifying

complicated mathematical formulae. *MACSYMA*'s knowledge base embodies the knowledge of several expert mathematicians.

PROSPECTOR [3] is a highly rated system designed to assist geologists in evaluating findings for mineral exploration. The work done on the *PROSPECTOR* system has helped to identify other related application areas including oil well log analyses.

One of the first expert systems to be used in a routine industrial setting is *RI* [5]. Recently renamed *XSEL*, this system is an expert at configuring Digital Equipment Corporation VAX-11 computing systems from customer order specifications. DEC has claimed that *XSEL* has saved considerable effort in its day to day operations dealing with the VAX line of computers.

A number of programs have been implemented for various applications in medicine. The best known example is *MYCIN* [10]. *MYCIN* was trained by a number of medical practitioners to provide advice on diagnosis, as well as therapy recommendations, for bacterial infectious diseases. In a related field, the *MOLGEN* [11] system is widely known for its ability to plan laboratory experiments for producing genetic material. Much of *MOLGEN* is presently being exploited for commercial purposes.

Each of these examples is interesting in their own right. However, the wide range of domains each has been applied to would seem to have little to do with CAD/CAM problems in general. In the following, we note certain analogies between various CAD/CAM problems with the application domains identified above. Thus, the methods employed in these expert systems may provide guidance for implementing solutions to CAD/CAM problems.

5. CAD/CAM problems

Probably the most advanced, and hence successful, CAD/CAM systems in existence today appear in the electronics industry. Systems have been implemented which can turnaround a logic design into a functioning integrated circuit in weeks rather than months. These systems have not been implemented with AI techniques. So why should the field of CAD/CAM be interested in expert systems?

The reason seems simply to be the complexity of the CAD/CAM tasks which demands more intelligence of its computational resources. In integrated circuit fabrication, the number of processes involved to implement functioning two dimensional circuits, constructed from a few possible materials is far smaller than the myriad of machining processes possible for three dimensional objects constructed from a wide range of materials. Choosing the proper processes and materials requires a considerable amount of intelligence.

Ideally, in CAD/CAM a design is generated using an engineering analysis such as a finite element technique, and the database used to generate that design is also common to the associated Computer Aided Manufacturing procedure. In fact, both the CAD and CAM portions still include a multitude of problem areas in both mechanical and electronic applications. AI has the prospect of impacting constructively in each of the CAD and CAM portions as well as in the manner in which they are connected via a common database. With respect to CAD aspects, speculations have been made about using AI in concert with finite element procedures. As a result, a potentially much greater number of engineers could be utilizing finite element procedures in a more consistent manner with many more application possibilities.

Automating design and manufacture of products, say in the automobile industry, poses an enormous range of problems and challenges fortuitously not generally encountered in electronics. From product conception and design, process and production planning and automated manufacture, the field of CAD/CAM has enough to keep it busy for a number of years to come. We will focus on a small set of problems, identified in [7], that present commercial CAD/CAM systems attempt to automate to some degree. (In fact, much of what we will be discussing here emphasizes the Computer Aided Manufacturing portion of the CAD/CAM problem.) These include:

- *Stock selection* - Using an inventory data base system, select the most appropriate stock item from a specification of a desired part to be machined.
- *Process planning* - Determining a plan of action for machining a specific part driven by data bases of geometric models of various parts
- *Production planning* - Integrating and scheduling a myriad of process plans for a factory to most cost-effectively produce the desired parts.
- *Production monitoring* - Verifying that the various factory workstations are correctly machining the desired parts.

To maintain our truth in advertising, the reader should be warned that the author is by no means an expert in CAD/CAM, nor of the general manufacture process. Thus, it is with a somewhat naive view that we offer the following speculation.

5.1. Stock Selection

Existing CAD/CAM systems rely on standardized data bases of stock parts or families of parts, with associated process plans for their production. Users of such systems attempting to plan the machining of a new part first select the most appropriate part family along with its associated process plan. Appropriate part families, though, are typically selected by simple codes as specified by the user. Thus, part family selection is left primarily to the ingenuity of a skilled human expert.

A more automated approach to this aspect of CAD/CAM might include the ability for an intelligent system to select its own candidate part families from a high-level specification of the desired part. This selection process seems ripe for expert systems technology.

Stock selection has a close parallel to candidate molecule selection implemented in the *DENDRAL* expert system. (A similar analogy may be drawn between stock selection and mineral selection as in the *PROSPECTOR* system.) From very general specifications, or constraints, imposed by preliminary NMR spectral data, *DENDRAL* successfully restricts its attention to a relatively small number of candidate molecules. To be sure, the problem *DENDRAL* attempts to solve requires enormous computational resources. One of the very useful approaches that was implemented in *DENDRAL* was a planning mechanism that redefined the original problem to a much simpler problem. The NMR spectral data was used for the redefinition process.

The selected molecules, or molecular families, are then subsequently refined

and detailed by other intelligent processes within *DENDRAL* to eventually produce an accurate description of an unknown chemical compound under investigation. The analogy of three dimensional molecules to three dimensional geometric parts is a compelling one. Thus, the methods employed in *DENDRAL* seem to provide guidance in the implementation of similar processes for part family selection.

Furthermore, the geometric and mathematical models used to represent stock items may be manipulated by systems such as *MACSYMA* to produce simplified and perhaps more intelligent indexing schemes. The generalization-based (hierarchically organized) indexing schemes typical of many AI expert systems seem a perfect choice for representing information about the relationships between various parts and families of parts. Wasserman and Lebowitz [13], for example, have been studying such an approach for an AI system which deals with patent abstracts of disc storage devices.

5.2. Process Planning

Once a particular stock item and part family is selected the next phase to be considered is process planning. In this case an existing plan, extracted from a data base, for machining various geometric surfaces is modified and verified to construct the desired part.

The *RI* system draws the closest parallel in this case. *RI* has great expertise in configuring VAX computing systems. Although the VAX computer is a complicated system of integrated parts, the number of peripheral devices and options available for a VAX are not as large as say the number of possible integrated circuits available to implement a complete VAX system. Thus, the configuration of a VAX is rather tightly constrained by DEC's available device offerings. Thus, the analogy here lies between the tight constraints of a VAX configuration with the constraints imposed by the canonical process plans.

RI is implemented by a rather simple problem-solving process similar to that described earlier for *ACE*. The general problem of VAX configuration can be well enough defined by a static sequence of subproblems, each to be solved in the same order each time guaranteeing an appropriate solution. For process planning, it would seem that similar problem constraints exist for each part family. For example, a surface should be scheduled for drilling before the introduction of bolts and nuts is planned.

5.3. Production Planning

Similar methods may be used for production planning. After a number of detailed process plans have been constructed, the next phase is integration and scheduling of the various plans for implementation in a factory with a number of machining workstations. In some sense this is a more difficult problem to solve automatically than process planning since many more possibilities exist. The process plans to be implemented have a number of available workstations to be scheduled, optimizing on various production parameters such as part trajectories and conveyor speeds. Thus, the combinatorics of this problem seem more dramatic than that for process planning.

This particular CAD/CAM process has obvious connection to a general problem studied for many years by AI researchers planning. The *MOLGEN* expert system is one of the best known examples of an AI planning system.

As noted, *MOLGEN* assists molecular geneticists in planning laboratory procedures to manufacture specific genetic material. Beginning with general skeletal plans in a very abstract setting, the system repeatedly refines its laboratory plan using a great deal of knowledge about molecular genetics, until a final detailed plan is produced.

One of the key approaches used in *MOLGEN* is constraint propagation. As the system chooses specific objects and processes to replace the more general abstract specifications, constraints are imposed on the entire slowly forming plan. The program is very adept at verifying that no newly introduced object will violate existing constraints in the plan specification. If so, *MOLGEN* makes different choices opportunistically to maintain the consistency of the plan. The final detailed plan produced by the system is guaranteed in this way not to violate any constraints as specified by the initial proposed problem.

The great variety of processes available, as well as initial laboratory material to be used has great similarity to the problems encountered in production planning. Here the great variety of process plans for machining three dimensional geometries can be expected to lead to an enormous number of possible schedules of machining processes -

6.4. Production Monitoring

After stock is selected, process plans constructed, the factory scheduled, we have finally come to production monitoring. In this case, our idealized CAD/CAM system would verify that the various factory workstations are correctly machining the desired parts on time. In the event of a failure of the manufacturing system, the specific workstations at fault must be identified and diagnosed.

It is a rather obvious suggestion to identify the parallel between an automated factory system and the human body. *MYCIN*, as noted, is designed as a consultant to medical practitioners to diagnose infectious diseases. Beginning with a general set of hypotheses concerning the ill patient, *MYCIN* asks specific questions including requests for the results of laboratory tests to elaborate the most likely causes and sites of infection, as well as to recommend the most appropriate therapy. Such a diagnostic tool would be invaluable for a CAD/CAM monitoring facility for an automated factory.

Presently, several CAD/CAM systems provide data bases of numerically controlled (N/C) tool programs. Each such program includes not only geometric models of parts but N/C tool instructions as well. We can draw a close analogy between N/C tools to human organs, and N/C tool instructions to organ function.

Lest we paint an overly simplified picture of the manufacturing process, we note the huge complexity involved in monitoring all aspects of the manufacturing process. This complexity is a result of the large complexity of its constituent processes and their interaction. However, the human body is no less complex a system, and in some sense is a more challenging device. Modern medicine unfortunately has a rather incomplete model of how our bodies function. Nevertheless, medical practitioners are adept at diagnosing and treating a myriad of diseases using sophisticated testing devices. The automated factory would seem to provide a better, or at least a more tractable problem environment since a precise model of the processes within exists. In our opinion, therefore, the methods employed in *MYCIN* seem to have sufficient power to have direct applicability to the problem of monitoring and diagnosing faults in an automated factory.

5.5. Use of Data Bases

Finally, how does *ACE* fit in to all of this speculation? The reader will note that several of the identified CAD/CAM processes are typically driven by (conventional) data base systems. Each such data base is fine-tuned to adequately assist human users to effectively solve the various problems at hand. It is our final assertion that the methods employed in *ACE* provide strong evidence that expert system technology can be brought to bear on various CAD/CAM domains. Indeed, our experiences in implementing the *ACE* knowledge base convinces us that in a very general way data base applications in existence today are ripe for this technology.

It is often written in the literature on expert systems that 5 man-years of effort is required to implement an effective knowledge base. The developers of *ACE*, on the other hand, completed the knowledge base in 8 man-months. There are many reasons for such a dramatic difference in implementation time. The existence of the data base helped to solve one of the most difficult aspects of building an expert system. The appropriate information, as well as methods for its manipulation and retrieval, for the problem at hand had been implemented, tested and debugged, ready for use by humans as well as computer programs. Thus, the knowledge representation problem had been greatly simplified. This saved an enormous amount of time for the system developers who concentrated on the more important issue of transferring problem-solving expertise from human expert to computer. It is our opinion that the same cost savings will appear when considering similar approaches to merging AI knowledge bases with CAD/CAM data bases.

6. Conclusion

The development and introduction of a new technology into industry generally comes with a rather high price tag. CAD/CAM is no exception. However, without much more of an investment of resources, present CAD/CAM systems can be improved dramatically with the introduction of expert systems augmenting their present capabilities. With skilled knowledge engineers, it does not appear that the addition of knowledge-based programs will be prohibitively expensive. Indeed, our conclusion is that present AI techniques can in the near future be brought to bear to cost-effectively implement what CAD/CAM promises.

The successful implementation of AI techniques in the CAD/CAM area will require the cooperative interactions between AI scientists and engineering analysts as well as manufacturing engineers, all attempting to build bridges across each other's disciplines.

Acknowledgments

We wish to thank Nicholas Perrone whose careful reading of an earlier draft of this paper substantially improved its presentation.

REFERENCES

- [1] Buchanan, B. G. and Feigenbaum, E. A. *DENDRAL and Meta-DENDRAL. Their applications dimension. Artificial Intelligence*, 11:5-24, 1978.
- [2] Davis, R. *Applications of meta-level knowledge to the construction, maintenance and use of large knowledge bases*. Rep. No. STAN-CS-76-552, Computer Science Department, Stanford University, 1976.
- [3] Duda, R., Gashng, J and Hart, P E Model design in the *PROSPECTOR* consultant system for mineral exploration. In D Michie (Ed.), *Expert systems in the micro-electronic age*, Edinburgh University Press, 153-167, 1979.
- [4] Forgy, C and McDermott, J, "OPS, A Domain-Independent Production System Language". *IJCAI* 5, 933-939, 1977
- [5] McDermott, J, "RI. The Formative Years". *AI Magazine* 2 21-29, 1981.
- [6] Moses, J. "A MACSYMA Primer". Computer Science Lab, Memo 2, Massachusetts Institute of Technology, 1975
- [7] Nau, D. "Issues in Spatial Reasoning and Representation for Automated Process Planning". Proc Workshop on Representation and Processing of Spatial Knowledge, 1983
- [8] Newell, A. "Production Systems Models of Control Structures". In W Chase (editor), *Visual Information Processing*, Academic Press, 1973
- [9] Rychener, M D. *Production Systems As a Programming Language for Artificial Intelligence Applications* Computer Science Department Carnegie-Mellon University 1975
- [10] Shortliffe, E H. *Computer-based medical consultations: MYCIN*. New York American Elsevier 1976
- [11] Stefik, M. "Planning with constraints". Doctoral dissertation, Computer Science Department, Stanford University, 1980
- [12] Stolfo, S. J and Vesonder, G T *ACE An Expert System Supporting Analysis and Management Decision Making* *Bell System Technical Journal* (to appear)
- [13] Wasserman, K and Lebowitz, M. Representing complex physical objects' *Cognition and Brain Theory*, 5 3 1983