Industrial Applications of Artificial Intelligence

Mark S. Fox
Intelligent Systems Laboratory, Robotics Institute, Carnegie-Mellon University, Pittsburgh, PA 15213, USA

This paper reviews current and future applications of Artificial Intelligence (AI) and Knowledge-Based systems to manufacturing. This is not a review of robotics technology, but focuses, instead, on manufacturing decision problems. Manufacturing, in this case, refers to the entire product life cycle: product design, production planning, production, distribution, and field service and reclamation. The review focuses on where, at each point in the product life cycle, there are problems to be solved, where AI is currently being applied, and where it may be applied in the future. The results of a recent survey note that research and development in this area has increased significantly in the 1980s. Most focus on spot applications of the technology. More recent work is taking a systemic view of manufacturing.

Keywords: Artificial Intelligence and Manufacturing, Knowledge-Based Systems, Computer-Aided Design, Planning, Scheduling, Diagnosis, Computer Integrated Manufacturing.

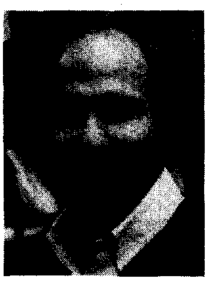

Mark Fox heads the Intelligent Systems Laboratory of the Robotics Institute, at Carnegie Mellon University. He is also a founder and Chief Scientist of Carnegie Group Inc. His primary interests are: Artificial Intelligence and Knowledge Based Management and Manufacturing Systems. At present, his laboratory is extending artificial intelligence techniques to the design and construction of engineering, production control and management systems for "flexible" factory organizations. Dr. Fox received a B.Sc. in Computer Science from the University of Toronto, and Ph. D. in Artificial Intelligence from Carnegie Mellon University.

North-Holland
Robotics 2 (1986) 301–311

1. Introduction

The purpose of this paper is to review current and future applications of Artificial Intelligence (AI) and Knowledge-Based systems to manufacturing. This is not a review of robotics technology, but focuses, instead, on manufacturing decision problems. Manufacturing, in this case, refers to the entire product life cycle:

- product design,
- production planning,
- production,
- distribution, and
- field service and reclamation.

The review focuses on where, at each point in the product life cycle, there are problems to be solved, where AI is currently being applied, and where it may be applied in the future.

2. What is Artificial Intelligence?

Artificial Intelligence is the science concerned with the creation of machine intelligence which is able to perform tasks heretofore only performed by people. Much of this machine intelligence is symbolic and heuristic. Artificial Intelligence deviated from the juggernaut of computer research in the early 50s by exploring how computers can be used for more than just numeric processing. Back during the days when languages like COBOL and FORTRAN were being defined, people at Carnegie-Mellon University and MIT were investigating the simulation of human problem-solving on a computer. Some of the first programs at that time were being applied to solving logic problems as found in the book "*Principles of Mathematics*" by Whitehead and Russell [23]. Problems such as chess, checkers, image understanding, etc. began to be investigated. As a matter of fact, one interesting set of problems that people chose to demonstrate the AI techniques could solve problems at the same level that humans could, were chosen from the intelligence test that we normally give to students to measure their IQ [6]. It so happens that computers are very good at solving them. The difficulty in the development of machine intelligence lies in programming computers to perform common sense reasoning, i.e. reasoning about every-day occurrences which people find easy to do.

The types of problems which AI attempts to solve are non-linear and combinatorially complex (e.g. planning/scheduling, image understanding, etc.). The impact of their being non-linear is that there do not exist algorithms which will provide optimal solutions in polynomial time. Hence, the use of symbolic, heuristic knowledge, or "rules of thumb", play a major role in AI systems.

AI research can be divided into two basic categories. Knowledge representation is concerned with *how* to represent knowledge in a computer understandable form, so that systems can act in an intelligent manner. Consider a description of an activity that occurs on a factory floor.

> The milling operation precedes inspection. It is composed of two steps: setup and run. Setup takes an hour and piece time is 30 minutes. Two resources are required by the operation. A wrench and an operator. The wrench can be found in the tool crib and is only required during setup. The operation is performed in cost center 84.

The question is: How would one represent the knowledge embedded in that paragraph in a computer? Typically, a record in a database which contains fields in some language is created. Describing the basics of an operation, there may be a field in a record which describes the tooling required, type of operator, who the operator is, setup time, run time, and the next operation in a sequence of operations to produce a product. The problem is: "How does a computer understand the meaning of the record?" The answer is that the program which uses it projects its own interpretation. The goal of knowledge representation research is to move away from ad hoc representations of knowledge to a semantics-based representation which identifies the levels of representation and their machine interpretation [2]. One level of representation, the conceptual level, provides a standard semantics which can be used across organizations and tasks, such as simulation, scheduling, accounting, etc. Fig. 1 is an example of the representation of the knowledge in that paragraph. It is relational in form. Knowledge representation research focuses on the identification

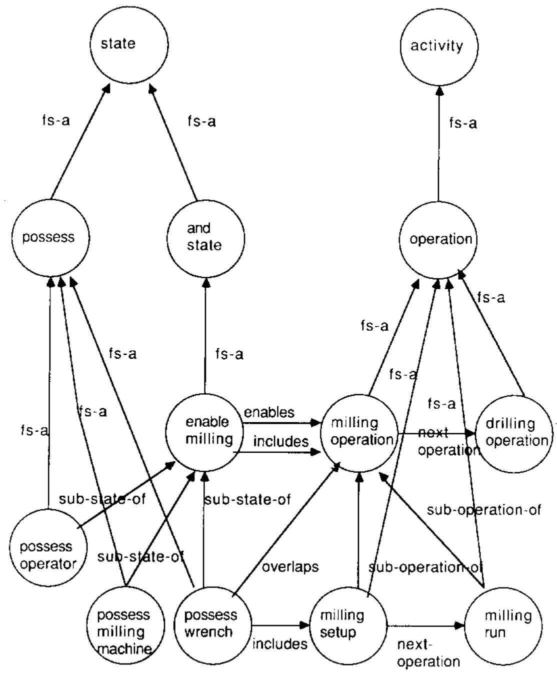

Fig. 1. Activity semantic network.

standard types of nodes and relations. Fig. 1 divides the knowledge into two types: activity and state. A state description describes a snapshot of the world before an activity is performed. For example "cost center 84 possesses a wrench" is a state description. It must be true in order to enable the milling activity to occur. States and activities are linked via causal relations. A state describes what must be true of the world to enable an activity to occur. In addition, activities may be defined at multiple levels of abstraction. Milling is refined into two sub-activities: the setup and running of the machine. Lastly, we must represent time. Setup, in time, occurs before the milling run. Time is not absolute, it is relative. When describing the factory floor, it is atypical to use absolute time periods, instead, activities are described as preceding each other; and hence, once the time of one activity is determined the time of all the other activities related to it can be determined.

This is a small glimpse of what people in AI have been doing in the area of knowledge representation. For more details on conceptual level representations of activities see [26].

The second category of AI is research. Problems are solved by performing search. Search in scheduling may be described as a search which incrementally builds alternative schedules. First, starting with a single order, it generates alternative first operations. Then, for each operation, it generates alternative machines on which to perform that operation, and for each machine it generates alternative queue positions, that is, times to perform that operation. There may be other alternatives such as alternative shifts and substitute tool-

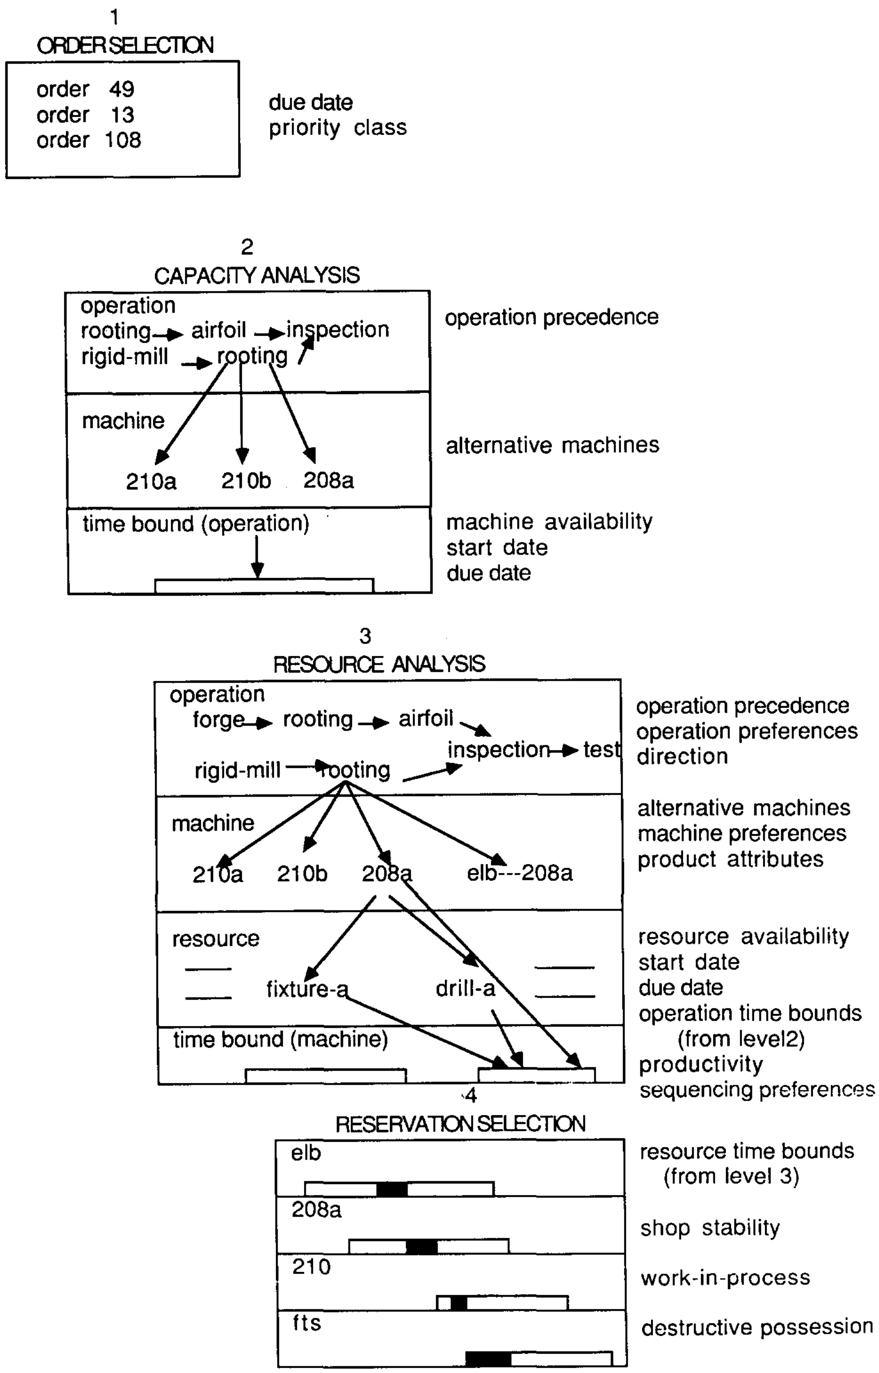

Fig. 2. ISIS search hierarchy.

ing and fixtures, which expand the tree into a larger network. Once the first operation is fully defined, the search proceeds with the next operation that follows it. Search can be performed in a forward manner or in a backward manner, starting from the due date.

This is a very simple example of search. This type of search is similar to playing chess by considering all the alternatives. In a single factory which has 85 orders, 10 operations and only one substitutable machine, there are over 10^{800} alternative schedules that one could create. So, while this kind of search is an interesting, theoretical technique, practically, it is unusable because there are too many alternatives to consider. So, search has to be much smarter; in fact, search has to be structured in a manner that the search space is reduced from 10^{800} to something much smaller and manageable. The requirement to reduce the combinatorics of search has led to more complex and sophisticated search architectures. Fig. 2 depicts the search architecture for the ISIS job shop scheduling system [10]. ISIS has a hierarchical search architecture where it divides search at each level into three phases:

1. pre-analysis of the search space to choose the operators which generate the search space,
2. the search phase which performs the actual problem-solving, and
3. post-analysis which analyzes the results of the search to identify whether it was successful or not and whether it should be performed again with a change in operators and constraints.

These phases are performed at each level within the search hierarchy. The first level chooses what order to schedule, the second level performs a search where it considers only a subset of the information in the factory, such as machine capacity, alternative operations, due dates and start dates, and tries to understand what the capacity is that is available, identify bottlenecks etc. The third level considers more information about the search space, substitutability of machines, tooling, shifts, constraints on how you perform your manufacturing, working in process, time constraints, how important it is to meet due dates, quality constraints, cost constraints, etc. The search results of higher levels constrain the processing of lower levels. Finally, the fourth level determines the time periods during which resources are allocated to operation. So, the system successively refines its schedules by starting from the aggregate and refining it via search to a detailed complete schedule by including more and more detail. Search-based systems like ISIS use sophisticated search techniques to reduce the size of the search space while still generating a good schedule.

The combination of representation and search are, in essence, the two legs of AI. It is upon these two legs that knowledge-based systems are constructed.

3. Issues in Manufacturing

There are a number of reasons why Artificial Intelligence may be important in the manufacturing environment. Consider the following issues. The scarcity of expertise is endemic to many corporations. Each time I visit a manufacturing organization there is at least one request of the following type: "Mark, we have an expert in a particular area and he/she is going to retire in two years. We want to save that knowledge for the organization." Not only do they want to capture that knowledge before the person retires, but they also want to distribute it throughout the organization. The issue of capturing scarce expertise and distributing it is a major problem faced by manufacturing organizations today. It occurs in tasks such as process planning, diagnosis of machinery, scheduling, and engineering design.

Another issue being faced in manufacturing is decision complexity. Decision complexity arises when there is a large number of choices from which to choose. They may be engineering choices of how to design a product, or they may be scheduling choices of how to produce the product. It is the case that flexible manufacturing systems are exacerbating the problem. The added flexibility on the factory floor provides the scheduling person with more alternatives, hence more choices of how to produce product. Today, making scheduling decisions in very rigid manufacturing environments is too difficult. As these environments become more flexible, the complexity of decisions increases. The same is true in engineering: the product complexity is increasing and, hence, design complexity increases.

Information is also becoming more complex. Many people use the term 'paperless factory' or 'paperless office'. The problem of how to get information from the factory floor on-line is now being solved. This creates another problem: how to reduce on-line information to only what is necessary for an individual to make a decision? Putting information on-line does not make it accessible to the individual decision maker. Reducing 1000 pages to a single screen full of pertinent information is an unsolved problem today. The intelligent reduction of information on an individual-by-individual basis is a long way off.

Decision timeliness is another issue which is coupled with decision complexity. Not only is decision complexity increasing, but the time to make a decision is decreasing. For example, programmable automation reduces the setup time of machinery, hence the scheduling decisions have to consider more alternatives in less time: how can faster, smarter decision-making systems be built?

Lastly, there is an issue of coordination. It is now known that design is intimately connected with production, distribution and field service. If a design is not optimized for manufacturability, assembly, distribution, or field service, then it will increase its cost of manufacturing with the possibility of reduced quality. The question is: how can designs be coordinated with all the down-stream activities?

These are some of the issues that affect the ability to increase the quality and the productivity of manufacturing operations. The question arises again as to whether Artificial Intelligence is a useful technology.

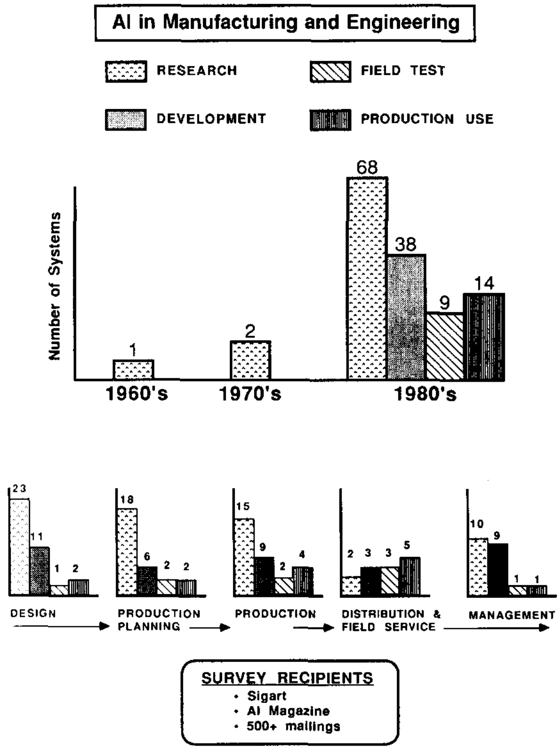

Fig. 3. AI in manufacturing survey.

4. AI in Manufacturing Survey

The results of a recent survey show the extent to which AI is being applied to manufacturing problems. The survey shows that in the 60s and 70s there was very little work being done (Fig. 3). But in the 1980s there are at least 68 systems in research, 38 in development, 9 in field test and 14 in production use. With a response of about 125 systems this represents about a quarter of the number of real systems being investigated today. There are closer to 500 systems that are under construction around the world today using Artificial Intelligence techniques.

The survey demonstrates that a number of people believe that AI will have an impact. The question is: where? In the following, I review the areas of impact by examining each phase of manufacturing.

5. AI in Design

The first phase of the product life cycle is design (Fig. 4). Design is composed of a product specification followed by the actual design and then its validation. In parallel, a number of management activities are performed. AI is being applied to each of these phases. Today, specifications are natural language descriptions written, in the case of commercial products, on a few pages, or in the case of military products, on thousands of pages. It is necessary to analyze these specifications to determine their completeness and consistency.

Artificial Intelligence has been applied to the acquisition and analysis of specifications. XSEL [20] is an example of a system which works with the computer sales person to acquire customer product requirements. The XSEL system has a natural language interface for the specification of

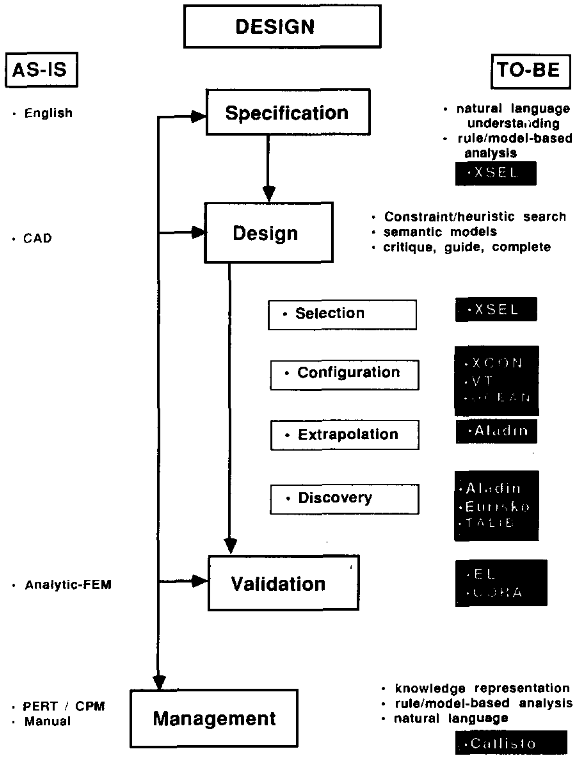

Fig. 4. AI in design.

customer goals; it analyzes customer needs and queries the customer in order to complete the model of the customer's needs. This system is in production use. XSEL deals with a portion of the specification problem. The general ability to analyze complete specifications, identifying consistencies and incompleteness is beyond the state of the art but is something to which AI techniques can be applied.

Once the specification is received and its completeness and consistency verified, design may proceed. Design is one of the most difficult and creative of the manufacturing tasks. It requires not only education but much experience. Consequently it is a knowledge-intensive task which utilizes both analytic and heuristic knowledge, and therefore makes AI techniques appropriate.

There are at least categories of design. The first is *selection*, which maps functional requirements onto the attributes of an existing product line, hence selecting a product from an existing set. For example, XSEL maps user functional requirements onto product attributes and selects the appropriate products.

Configuration is a second type of design and is a little more difficult. The problem is to combine semi-finished components to form a final product. There are a number of successful configuration systems. XCON/R1 [19] configures the VAX and PDP/11 lines of computers. It is in production use today. VT [18] is another system which configures elevators and is in field test. OCEAN [24] is another computer configuration system. The XCON system has been in use for three years. It has configured over 100.000 computer orders at Digital Equipment Corp. and is, on the average, over 99% correct, which exceeds the ability of human configurers. It is also faster, configuring systems in 3–5 minutes versus one hour to three hours for human configurers.

Extrapolation is the third type of design. Here, an existing product is altered to meet the customer's specifications. ALADIN [7] is an example of an extrapolation system. It takes the specification of the properties for an aluminum alloy, and alters an existing alloy's composition and thermal-mechanical processing, so that a new alloy which meets the customer's needs is defined.

Discovery is the most difficult type of design. It is similar to configuration in that it takes components and combine them. The difference is that configuration uses more semi-finished components, i.e., memory boards, CPU, busses, etc. for computers, whereas discovery begins with transistors and resistors. The distance in the functionality of the components from the end product determines whether it is discovery or configuration. There are systems which demonstrate the effectiveness of AI techniques on discovery problems. ALADIN can start just with aluminum to design an aluminum alloy. EURISCO [16] is a discovery system which has been applied to the design of VLSI circuitry. TALIB [14] is a system which designs electronic circuitry. All of these systems have been able to design/discover interesting functional systems out of basic primitives.

Once the design is complete, *validation* is used to verify that the form of the design behaves as defined by the functional specification. Simulation is used most often; for example, finite element analysis is used to test stresses in mechanical parts, and in electronic parts simulation is used to verify circuit behavior. The use of AI for validation differs in that it attempts to replicate what humans do when they validate a design: people have the ability to look at the structure of a design and infer functionality with minimal use of simulation. Two systems have been investigated in this area: the CONSTRAINTS system [30] infers functionality from the form of an electronic circuitry, and the CORA system [33] verifies relay protection systems which protect power lines.

Lastly, AI is also being applied to the management of engineering projects. In particular, the CALLISTO system [27] focuses on the management of product definition, and the activities performed to design the product. It provides a knowledge representation for the representation of dynamic objects and activities, and captures expertise for the management of activities.

6. AI in Production Planning

Production planning (Fig. 5) is the second stage of the manufacturing product life cycle. It takes a product design and definition of the production facilities as input. It forecasts customer demand, plans the process, lays out the facilities, specifies maintenance and trains the workers.

AI is being applied to each of the tasks. In the case of forecasting there are researchers in marketing groups who are looking at rule-based techniques for simulating market behavior. These systems go beyond the typical techniques which tend to predict the future based on the past by having a greater understanding of the individual's needs. The ROME system [15] focuses on another part of the forecasting problem, the analysis of resource plans. It is an intelligent VISICALC tm where the input is natural language and the system automatically analyzes the data to see if production goals are being met. If not, it figures out why.

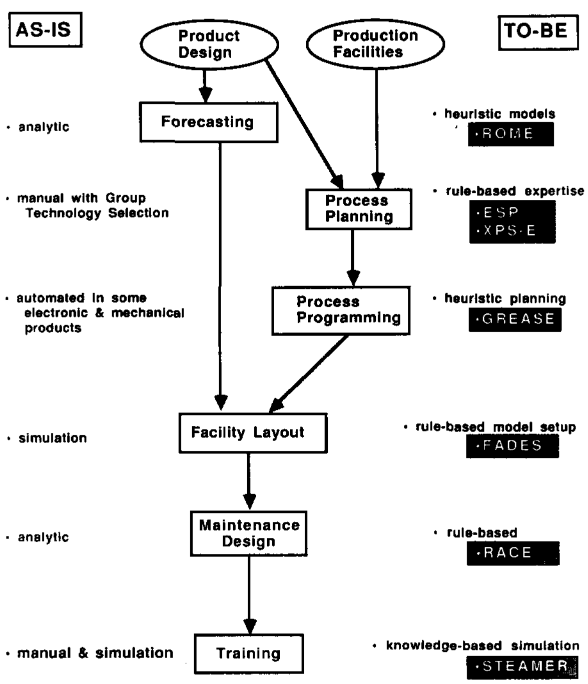

Fig. 5. AI in production planning.

A number of process planning systems are being investigated; ESP [17] has been developed for planning sheet metal production. XPSE, a continuation of the GARI system [5], is looking at the process planning of three-dimensional mechanical parts. Both are rule-based systems.

Process programming, for electronic applications programming, is straightforward and can be automated using conventional techniques. But for mechanical processes, such as the assembly of a carburator, we do not know how to automate the programming of the assembly process. The application of AI to this problem is being investigated for the assembly of a power supply [3].

Another aspect of the process programming task is the selection of supporting resources such as cutting fluids for machining operations. GREASE [22] is a system which selects cutting fluids for machining operations. It is now in field test.

Facility layout is concerned with how to organize production facilities based on process plans and product forecasts. The FADES system [8] has been used to select the parameters of an Operations Research technique, which then determines the facility layout.

In maintenance design, AI is being used to develop maintenance procedure. RACE [13] is one such system now in field test.

Lastly, training is an important aspect of manufacturing. The impact of poorly trained workers can be devastating on production. Typically, training is performed manually, though simulations have been used for aircraft and nuclear power stations. What AI systems bring to the table is the ability to capture human expertise and using it to train new workers. It is also the case that expert simulations using Artificial Intelligence are being used. For example, Simulation Craft tm [4] is an expert manufacturing simulator. It captures and provides scarce simulation, statistics and manufacturing expertise is support of the simulation life cycle: model specification, experiment design, and result analysis.

7. AI in Production

Production (Fig. 6) is concerned with the planning, scheduling, management of the shop floor, control of cells, inspection of product and maintenance of processes.

Scheduling is a combinatorial problem whose complexity, in many cases, exceeds the ability of human schedules. AI has been applied to scheduling. ISIS [10] is a system which uses constraint-directed search techniques for the scheduling of job-shops. Constraint-directed search uses constraints to reduce the combinatorics of the search resulting in schedules which satisfy some constraints while relaxing others.

IMACS [12] focuses on flow shop scheduling. The important concept in IMACS is its approach to managing a schedule. IMACS believes in Murphy's Law. Murphy's Law states that whatever can go wrong will go wrong. Rather than design the shop floor control systems to expect the schedule will be followed, it expects that deviations will occur and focuses on their detection and repair.

The same concept of looking for deviations is also used at the cell level by the TRANSCELL system [1] and the SCD system [32] developed at Hitachi. These are rule-based systems which react to change and identify the next activity to perform.

For inspection, both rule-based and model-based techniques are used to diagnose computers and printed wire boards. IDT [29] is a system for testing computers that are assembled, and IPWBIS [28] is a system for inspecting inner layers of printed wire boards and identifying where in the production process the error was introduced.

8. AI in Distribution and Field Service

Distribution and field service (Fig. 7) are the final phases of the manufacturing product life cycle. Distribution begins with the design of the organization: where to manufacture, how much to manufacture, how much to inventory, what transportation routes to use, what type of transportation to use, etc. Other distribution tasks are order entry, product installation planning, and diagnosis

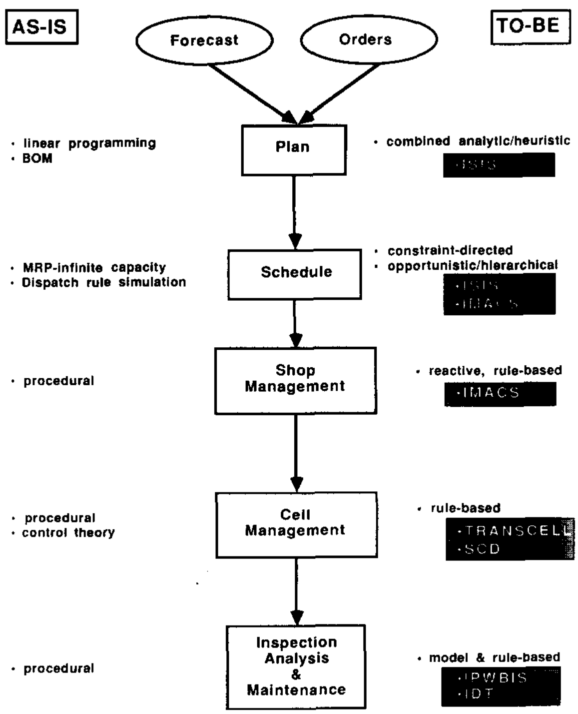

Fig. 6. AI in production.

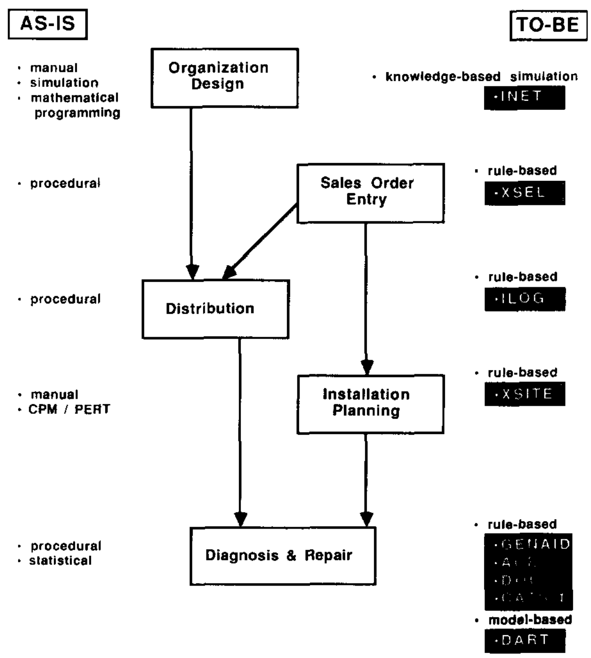

Fig. 7. AI in distribution and field service.

and repair of the product at the customer's site. AI again is being applied to each of these tasks. The INET system uses knowledge-based simulation techniques to model organization structures, simulate them and analyze automatically the results. XSEL is used for sales order entry. The output of XSEL is then provided to a distribution system, ILOG, which decides where to source the order, where to assemble, and how to ship it. Each of these systems uses rules to capture and apply expertise, and heuristic search to improve upon the solution.

Installation planning has also been attacked using AI. The XSITE system [21] decides how to lay out the computer at the customer's site.

Lastly, diagnosis and repair is an area that has received a great deal of attention from AI. PDS (aka GENAID) [9] is a system which, using sensor input, performs diagnosis of steam turbines and generators, using sensor input. It has been in production use for more than a year. PDS is interesting because it has to worry about sensor inputs. The problem with sensor inputs is that the sensors degrade in many situations more quickly than the process itself. Diagnosis in this environment must determine whether the sensors are incorrect or whether the process itself is incorrect. That makes the problem a lot more difficult.

ACE [31] is a system which performs diagnosis of cable problems for telephone cabling. DOC is a system which performs diagnosis of computers and CATS is a system for the diagnosis of diesel locomotives. These are all rule-based systems which use heuristics, knowledge of situation and action. Current research in AI focuses on model-based reasoning, representing the physical structure and operating theory of the product, i.e. chemistry, physics, thermo dynamics, etc. For example, the DART system [11] diagnoses computers based upon their structure.

9. Future Trends

In the preceding, I have examined each phase of manufacturing: design, production planning, production, distribution, field service; and I have identified point applications of Artificial Intelligence. The question is: What is the future of applying Artificial Intelligence to manufacturing? What needs to be done is to take a system's view of the manufacturing process. One view is strategic: How can a product be designed so that it optimizes all the down-stream activities? I call this 'Design Fusion' (Fig. 8). The problem is to represent and utilize all the knowledge about planning, production, distribution and field service, during the design of the product, so that down-stream activities are optimized.

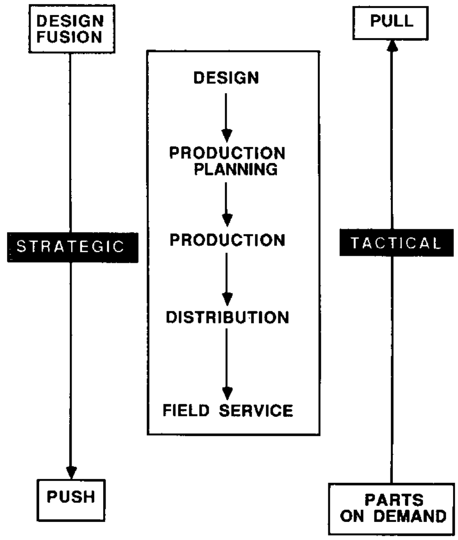

Fig. 8. Future trends.

The second system's view is called 'Parts on Demand': How is an existing product pulled through the organization? Parts on demand decisions are concerned with tactical issues of: Where to get the product? How to get it more quickly? If the product does not exist, how can it be produced quickly? Or in the case where the design of the product no longer exists, how can it be reverse engineered? By taking a more global strategic view and tactical view of pushing new designs and pulling existing products within the organization, it becomes possible to see how spot applications of AI can be further optimized.

10. Conclusion

In conclusion, there exist a number of applications of AI in manufacturing today. They are beginning to impact manufacturing both on the shop floor and in engineering design. My expectations are that the number of systems will continue to increase at an even larger rate as corporations acquire more AI expertise and feel more comfortable about its application. More systems which capture scarce expertise and make it available throughout the organization will be created. Systems that enhance our problem solving by making better decisions more quickly will be created. Systems that integrate more knowledge about the factory floor and, hence, make better decisions will be created. There will be increased accessibility to these systems by people who are not computer-oriented through the use of natural language and explanation facilities.

The future appears to be rosy, but there are barriers looming along the road. The problem is that there are few people who really understand AI. There are people who have book knowledge of AI but do not understand AI from a point of view of actually building real, large systems. It so happens that AI is like a guild. It is a guild in a sense that there are masters, journeymen and apprentices. The problem is, there are not enough masters for the apprentices to work with. Because of this, many organizations will not have enough masters or any masters at all to guarantee a project's success. There will be a number of failures over the next five years, as corporations stub their toes in applying the technology. The possibilities of success may be enhanced if the corporation works with masters, starts small, lowers its expectations, and makes sure that the problem can be done by people before solving it with a computer. Only then will it be possible to successfully apply Artificial Intelligence to industrial problems. AI has a lot of promise but it must be applied cautiously.

References

[1] Bourne, D., and Fox, M.S. (1984), "Autonomous Manufacturing: Automating the Job-Shop", *Computer*, Vol 17 No 9, pp. 76–88.

[2] Brachman, R.J. (1979), "On the Epistemological Status of Semantic Networks", in: *Associative Networks: Representation and Use of Knowledge by Computers*, N.v. Findler (Ed.), pp. 3–50, New York, Academic Press.

[3] Chang, K.H., and Wee, W.G. (1985), "A Knowledge Based Planning System for Mechanical Assembly Using Robots", *Proceedings of the 22nd IEEE Design Automation Conference*, pp. 330–336.

[4] CGI (1986), "*Simulation Craft*", Carnegie Group Inc., Commerce Court at Station Square, Pittsburgh, PA 15219.

[5] Descotte, Y., and Latombe, J.-C. (1981), "GARI: A Problem Solver That Plans How to Machine Mechanical Parts", *Proceedings of the Seventh International Joint Conference on Artificial Intelligence*, pp. 766–772, Vancouver, B.C.

[6] Evans, T.G. (1968), "A Program for the Solution of Geometric-Analogy Intelligence Test Questions", in: *Semantic Information Processing*, M. Minsky (Ed.), Cambridge, MA, The MIT Press.

[7] Farinacci, M.L., Fox, M.S., Hulthage, I., Rychener, M.D. (1986), "The Development of ALADIN, an Expert System for Aluminum Alloy Design", *Artificial Intelligence in Manufacturing*, Thomas Bernold (Ed.), Springer-Verlag, to appear.

[8] Fisher, E.L. (1984), "Knowledge-Based Facilities Design", Ph.D. Thesis, School of Industrial Engineering, Purdue University, W. Lafayette, Indiana.

[9] Fox, M.S., Lowenfeld, S., and Kleinosky, P. (1983), "Techniques for Sensor-Based Diagnosis", *Proceedings of the International Joint Conference on Artificial Intelligence*, August 1983, Los Altos, CA, William Kaufman Inc.

[10] Fox, M.S., and Smith, S. (1984) "ISIS: A Knowledge-Based System for Factory Scheduling", *International Journal of Expert Systems*, Vol 1, No 1.

[11] Genesereth, M.R. (1982), "Diagnosis Using Hierarchical Design Models", *Proceedings of the National Conference on Artificial Intelligence*, pp. 278–283, Los Altos, CA, William Kaufman Inc.

[12] Haley, P., Kowalski, J., McDermott, J., and McWhorter, R. (1983), "PTRANS: A Rule-Based Management Assistant", Technical Report, Computer Science Dept., Carnegie-Mellon University.

[13] Honeywell (1985), Personal communication with Honeywell Corporation.

[14] Kim, J., and McDermott, J. (1983), "TALIB: An IC Layout Design Assistant", *Proceedings of the National Conference on Artificial Intelligence*, pp. 197–201, Los Altos, CA, William Kaufman Inc.

[15] Kosy, D., and Dhar, V.S. (1983), "Knowledge-Based Support System for Long Range Planning", Technical Report, Robotics Institute, Carnegie-Mellon University, Pittsburgh, PA, December 1983.

[16] Lenat, D.B. (1983), "Theory Formation by Heuristic Search. The Nature of Heuristics II: Background and Examples", *Artificial Intelligence*, Vol 21, No 1, pp. 31–60.

[17] Lindsay, K.J. (1985), "Frame-Based Knowledge Representation for Processing Planning", *Proceedings of the First Annual Aerospace Applications of Artificial Intelligence*, Dayton OH, Wright-Patterson Air Force Base Materials Lab, September 1985.

[18] Marcus, S., Stout, J., and McDermott, J. (1986), "VT: An Expert Elevator Configurer", Technical Report, Computer Science Dept., Carnegie-Mellon University, Pittsburgh, PA, to appear.

[19] McDermott, J. (1982), "Rl: A Rule-Based Configurer of Computer Systems", *Artificial Intelligence*, Vol 19, No 1, pp. 39–88.

[20] McDermott, J. (1982), "XSEL: A Computer Sales Person's Assistant", in: *Machine Intelligence 10*, J.E. Hayes, D. Michie, and Y.H. Pao (Eds.), John Wiley & Sons.

[21] McDermott, J., and Steele, B. (1981), "Extending a Knowledge-Based System to Deal with Ad Hoc Constraints", *Proceedings of the Seventh International Joint*

Conference on Artificial Intelligence, Vancouver, B.C., August 1981.

[22] Mogush, J.E., Carrega, D., Spirtes, P., and Fox, M.S. (1986), "Issues in Diagnosis and Treatment: A Case Study in Cutting Fluid Selection", Technical Report, Robotics Institute, Carnegie-Mellon University, Pittsburgh, PA, to appear.

[23] Newell, A., and Simon, H.A. (1956), "The Logic Theory Machine: A Complex Information Processing System", *IRE Transactions on Information Theory*, Vol IT-2, No 3, pp. 61–79.

[24] NCR (1985), Personal communication with National Cash Register Corporation.

[25] Reddy, Y.V., Fox, M.S., Doyle, K., and Arnold, J. (1983), "INET: A Knowledge-Based Simulation Approach to Distribution Analysis", *Proceedings of the IEEE Computer Society Trends and Applications*, National Bureau of Standards, Washington, DC.

[26] Sathi, A., Fox, M.S., and Greenberg, M. (1985), 'Representation of Activity Knowledge for Project Management", *IEEE Transactions on Pattern Analysis and Machine Intelligence,* Vol PAMI-7, No 5, September 1985, pp. 531–552. Also appeared as Technical Report CMU-RI-TR-85-17, Robotics Institute, Carnegie-Mellon University.

[27] Sathi, A., Morton, T., and Roth, S. (1986), "Callisto: An Intelligent Project Management System", AI Magazine, to appear.

[28] Seto, J., Thibadeau, R., and Friedman, M. (1982), "Automatic Inspection for Printed Wiring", Technical Report CMU-RI-TR-82-16, Robotics Institute, Carnegie-Mellon University, Pittsburgh, PA.

[29] Shubin, H., and Ulrich, J.W. (1982), "IDT: An Intelligent Diagnostic Tool", *Proceedings of the National Conference on Artificial Intelligence*, August 1983, Los Altos, CA, William Kaufman Inc.

[30] Steele, G.L. (1980), "The Definition and Implementation of a Computer Programming Language Based on Constraints", (PhD Thesis), MIT Technical Report AI-TR-595, Cambridge, MA.

[31] Stolfo, A., and Vesonder, G.T. (1982), "ACE: An Expert System Supporting Analysis and Management Decision Making", Technical Report, Computer Science Dept., Columbia University, New York, NY.

[32] Tashiro, T., Komoda, N., Tsushima, I., and Masumoto, K. (1985), "Rule-Based Control of Factory Automation Systems. An Approach with a Production System", Technical Report, Systems Development Lab., Hitachi Ltd.

[33] Westinghouse (1985), Personal communication with Westinghouse Electric Corporation's Productivity and Quality Center.