MIT/LCS/TM-56

RESEARCH ON EXPERT SYSTEMS

G. Anthony Gorry

December 1974

MAC TECHNICAL MEMORANDUM

RESEARCH ON EXPERT SYSTEMS

G. Anthony Gorry

December 1974

This research was supported by the Department of Health, Education, and Welfare (Public Health Service) under Grant number 1 RO1 MB 00107-01.

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

PROJECT MAC

CAMBRIDGE MASSACHUSETTS 02139

RESEARCH ON EXPERT SYSTEMS

G. Anthony Gorry

December 17, 1974

Expertise: Supply and Demand

To cope with the increasing complexity of social organizations and social processes, various segments of society have a growing need for experts. For example, experts are needed to help individuals or groups to cope with problems of medicine, law, education and the interface between the citizen and the government, as well as to get automobiles repaired and houses maintained. Part of this increasing demand for expert services results from an explosion of knowledge and rapidity of change that makes it extremely difficult for one person to comprehend a subject without a significant personal investment. Further, I believe that society's expectations are rising. People have come to expect services that formerly were thought to be luxuries denied to certain segments of the population. As time passes, more and more people see these services as being their right, something society owes them. Consider, for example, the increasing impetus for free medical care or legal service. Unfortunately, in certain important areas in society, the supply of expertise is not expanding as rapidly as the demand.

The supply of expertise is inadequate for a number of reasons. First, it may be absolutely inadequate. Perhaps there just aren't enough doctors to go around; enough lawyers to service the needs of the population; a sufficient number of architects to do the designs we want (i). In certain areas, although on an absolute count there are enough experts to meet the needs of society, they are poorly distributed with respect to these needs. That is, for geographical, sociological, or economic

----

(i) In some fields of expertise, there has been a *de facto* limitation on growth by the experts themselves. Examples are certain trade unions and medicine.

reasons, people in various population groups are denied access to expertise. Some people haven't got the subway fare to get from one side of Boston to the other to make a call on a medical expert. They may not even know of his existence. In rural areas, it is very difficult to obtain physician's services. There are towns with populations of ten to fifteen thousand that can't get a doctor to practice there. It is not that there are not enough doctors. That may compound the problem, but more important is the structural problem resulting from the unwillingness of physicians to practice medicine in rural areas. This is but one example of the problems of distributing expertise through the conventional mechanisms that society has developed.

Expanding the Supply of Expertise

Growing Experts

In order to expand the supply of expertise of a particular kind, society "grows" experts. Select and plant the seed; nurture that seed very carefully; and at some point he or she blossoms into a full grown practitioner of the art, whatever the art may be. For example, we select someone from college, then we send that person to medical school, then to a hospital where he is trained further; then out into the world as an expert in some medical specialty. Other professional schools provide essentially the same service for society.

The problems with this process are many. In general it is a very time-consuming, resource-consuming enterprise. Despite sharply rising fees, students are not paying professional schools all that it costs to educate them. Society, therefore, must invest enormous additional amounts of money to maintain these conduits through which people pass to become experts. Another problem is that the expertise that's generated in that way is generally fragile. People who have spent years learning their particular field have a great deal of difficulty keeping abreast of developments in that field, and it is very difficult for them to maintain a high level of competence unless they stay fairly close to a center that is producing experts. A doctor, for example, who is a specialist will not be a specialist very long if he goes into a rural community to practice. He may be an excellent doctor in that he is a good person, a humanitarian who treats the general needs of the population, but he will not remain an expert. The problem of the decay of his knowledge from lack of practice is compounded by the fact that technology and science move on, leaving him where he was the day he left medical school.

The newly-minted expert perceives this problem, and the psychological and emotional investment made in becoming an expert creates an inertia to stay in an environment which is maximally supportive. For example, the medical expert wants to stay near a teaching hospital. Therefore, even if we could "grow" these experts, even if we were willing to pay for the training of these people, we would find that it was still difficult to solve the distribution problem with respect to these various needs in society. (i)

Factorization of Expertise

Although we accept the thesis that expertise is a very valuable commodity and in limited supply, need we accept the thesis that the expert is expert in every aspect of his work? For example, can't we study a particular medical expert to identify the cutting edge of that person's expertise? If so, we could eliminate from his practice all those aspects which could be performed by someone of lesser skill, lesser training. Let's factor his expertise, reserving for him those pieces which are uniquely his and delegating those pieces that are more or less mundane. Is it necessary

----

(i) One approach is to provide incentive payments for experts, to pay more for doctors to practice out in the country; to pay more for lawyers to go into the ghetto areas; etc. In general, incentive payments have had relatively little effect, however, because people who are sufficiently motivated to become experts are not motivated solely by financial considerations.

for our medical expert to take a patient's history? Can't we have someone else take the history, summarize the results, and give them to the expert for his particular insight? Such factoring would allow the effective amplification of the expert's abilities, and thus alleviate the supply problem.

Some success has been achieved by factoring in certain areas, but this success has been limited. Most often, we don't know what constitutes the cutting edge, the real essence of what the expert does, and so we cannot identify a piece to give to a technician, a piece for a college graduate, and those pieces reserved for the expert. In the medical example, can we delegate the history to the technician? The expert says no. He is loathe to give up pieces of his domain, his territory. If we ask him "What do you do that you must do and what do you do that you can delegate to someone else?" he will identify a very small piece that he can delegate. He will contend that expertise is woven through all of his activities. He can't rely on someone else to take a history because that person might miss the one clue that would give him his key insight. Regardless of our assessment of his opinion, the facts appear to support it; we have been unable to successfully factor expertise in many areas.

Artificial Expertise

Another approach is to create artificial expertise. Rather than trying to understand how the human expert solves the problem, we can attack the problem *de novo* to devise a method for solving the problem in question which may or may not have anything to do with the way a human solves the same problem.

Consider the experience in some industrial applications. In every large manufacturing operation, there used to be an indispensable expert who managed inventories. He had a "sense" for when items ought to be ordered, how to associate the stocks he saw in front of him with anticipated demands, how to react to rumors of strikes, etc. He was the information repository and processor for the company with respect to inventories, a real expert who saved companies large amounts of money. Those people, however, are for the most part, dinosaurs of industrial production. They were well adapted to a certain niche, but technology, the computer, mathematics and operations research produced artificial expertise which was superior. The niche has been taken over by the artificial expert, the computer-based inventory system.

So we might be able, in a certain area, to construct artificial experts which perform at the level of the human expert currently functioning in that area. There are areas where we have at least some hunch, some insight, that leads us to believe that this could be done, areas where, if we were to attempt to build an artificial expert, we might well be successful. This approach merits attention, but frankly, I believe it will prove of limited value in complex domains.

The Use of Media

Finally, we can try to alleviate the problem of lack of supply and maldistribution of expertise by disseminating, through media such as books and papers, what the expert believes is the essence of his expertise. For example, the doctor who is an acknowledged expert in an area writes down what he deems are the important principles that guide his problem-solving (i). There's no doubt that one of the reasons for the rise in publications is this attempt to disseminate expertise about various areas. The difficulty with this approach is rather obvious to those who have tried to cope with the scientific literature, and I think it is fairly easy to understand all of the deficiencies in that way of transmitting information. Because it is passive and not active, there are major problems in trying to apply it to a given area. It can be helpful but it is not generally the solution

- - - - - - - - - -

(i) Of course the lack of success in attempts to factor expertise suggests that the doctor (or any expert) may not be able to clearly state such principles.

to the problem, because a major problem is knowing what knowledge is relevant in any given situation.

In summary, I believe that society faces a shortage in expert systems. Human experts are in short supply; they are not properly distributed with respect to the needs of society; but the mechanisms that society has developed for maintaining supply are now inadequate.

Computer-Based Expert Systems

Hence, I want to turn to the question of whether we can build computer experts, computer systems which would replicate the performance of experts in particular areas. If you have been exposed to the "computer culture," it won't surprise you that I suggest such a possibility. The large memory of the computer, coupled with the industry of its computing engine make it a likely vehicle to deliver expertise. One can imagine a computer system being an expert in a given field, shall we say an acknowledged expert in a particular domain in medicine. It has wonderful qualities. First of all it can be mass-produced, either in fact or in principle through time sharing systems. Thus, we will eliminate the problem of distribution. (i)

The other problem that this will alleviate, if not eliminate completely, is the problem of intellectual obsolescence we currently face because the time to grow an expert is so long. By the time the expert is trained and out in the field, he may suffer from many weaknesses in his knowledge, because things have changed, people have gotten better insights, and the things he was told and worked so hard to remember when he was in school are either irrelevant or just plain wrong. The expert computer program, by contrast, is relatively insensitive to the time at which information has been added to it. Therefore we can add knowledge to the program and thereby instantly disseminate it.

(i) Until the artificial intelligence people have gone a little bit further, we won't have to negotiate with our computer programs as to whether they prefer to live in the country or the city!

Another interesting aspect of building a computer-based expert for an area such as medicine, is that we in large part would have explicated expertise. In order to build such a system we will have to know what constitutes expertise; therefore, the system itself in large part will represent a theory of expertise in this particular domain, one that can be poked and prodded with various experimental techniques. If we used cognitive simulation to achieve our goal, that is, if we have tried to replicate human expertise, we then have an explicit description of that expertise which can be fed into the educational process.

One of the failings in professional education is that no one really teaches you to be a professional. No one really knows how to teach you decision-making in medicine or design in architecture, etc. You are taught by example so you have the good taste to recognize that this is a "good" X and this is a "bad" one, whatever X may be. Whereas, if one had a computer-based theory of expertise in a particular domain then, complex as it might be, it would represent an explicit statement of what was believed to be a sufficient body of knowledge and an organization of that knowledge for expert performance in the field.

We have a utopian view, a view of a revolutionary way of distributing expertise. It is appealing, this vision of the computer-based expert system, the system embodying sufficient amount of knowledge about a given area to be properly called an expert. Unfortunately, in most areas, such a day is a long way off. There are major problems, both intrinsic and technological, which stand between us and the realization of this vision. Closing the gap that exists between our current state of understanding and this particular future is the goal of computer science research on expert systems. I want to consider briefly some of the basic problems that are the focus of that research.

FUNDAMENTAL PROBLEMS

> The mental features discoursed of as the analytical, are, in themselves, but little susceptible of analysis. We appreciate them only in their effects. We know of them, among other things, that they are always to their possessor, when inordinately possessed, a source of the liveliest enjoyment. As the strong man exults in his physical ability, delighting in such exercises as call his muscles into action, so glories the analyst in that moral activity which disentangles.
>
> Edgar Allan Poe, *The Murders in the Rue Morgue*

The Nature of Expertise

We do not understand expertise. My facile use of the word so far is misleading at best. We do not have a theory of expertise. In any area that's reasonably complicated and interesting, our theory of problem solving, of cognitive process, or of expert performance is very gross and probably quite naive. So our first order of business is to ask "Who are the experts?" "How does their performance differ from the performance of non-experts?" and "What is the basis for their superior performance?"

We can think of a number of hypotheses, but none as yet has been tested in any thorough way. We can say, as many people have said, that the important part of expertise is really just experience, having seen thousands of situations. For example, we could claim that the major portion of expertise in chess is simple practice.

Alternatively, we could claim that experts are born, not made. Somewhere in the genetic code of a chess champion is a knight and that "explains" his superiority. The truth is, I suppose, as Simon once observed. Problem-solving is like athletics. In order to do certain athletic activities well, one must have a certain minimum amount of native ability. Some people have a maximum amount, and therefore, with relatively little effort they can outperform the rest of us. Examples are particularly common in sprinting and short distance running. On the other hand, there are people with moderate amounts of "natural" athletic ability who through hard work have risen to the top. Similarly, to play good chess, one has to have a certain ability and a certain amount of practice. Whether one is born with sufficient ability to avoid almost all practice is at the moment unclear.

The "Depth" Of Expertise

As we are unclear about acquisition of expertise, so too are we uncertain as to its "depth." One hypothesis is that expertise is deep. Expertise is deduction with formal ways of thinking as opposed to simply gathering large amounts of facts and matching stored experience. The paragon of deep expertise is Sherlock Holmes or Dupin in "The Murders In The Rue Morgue," detectives, who through various complex processes of deduction and powerful thinking, get to the root of the matter. In this view, we, too, all have the facts before us; in fact, the detective story is in large part structured so the facts are all in front of you. The detective triumphs, not because he has seen a hundred more crimes, but because he is a superior logician. Regardless of whether it is innate or acquired with practice, expertise is thought to be deep deductively.

In a rather polar view, expertise is seen as being broad but quite shallow. Problem-solving, in this view, is in the main pattern recognition. The expert is superior in his performance, because he has seen far more problem situations, and hence he remembers more "shadings" of these situations. He can make sharper distinctions among situations, and as a result, he can devise more appropriate or efficient strategies.

Support for this view of expertise as being anecdotal can be drawn from the educational process as it is currently constituted. Much of the time spent there is devoted to examples. A large number of cases is presented to the student to sharpen his perception of similarities and differences among them. Little, if any attention is paid to the explication of a "deep" problem-solving process.

Although such evidence is by no means conclusive, it does suggest that prototypes and examples play a large role in expert performance.

What is expertise? When we say that we want to build a medical expert, a legal expert, what is the basis for that expertise? How is it organized and how is it applied to any given problem. It is easy to state those problems, it is incredibly difficult to solve them for an reasonably non-trivial domain, and so they are the first order of business for researchers in the field of expert systems.

The Consultancy Role

To make matters even more difficult, there are other problems before us. We propose to put expert systems at the service of society or some subset of society. Imagine an individual coming to such a system for consultation. If this system is to consult with the user on the user's problem, it must have some notion of what it is to have a problem and what it is to be helped with a problem. Only in the most simple case is it immediately clear what the problem is. For most interesting applications, when the user approaches a consultant, he may be unable to articulate his problem. When a person goes to see a lawyer, he seldom says "I need an interpretation of the Land Title Act of 1967." If that were the case, we would be well along toward building expert systems. The client may not even know whether the lawyer is an appropriate expert, much less know how to define the problem. Hence, in most domains, we need to understand better what it means to state a problem and what it means to be helped with a problem.

Bounding the Problem Domain

> *Mrs. Eloise Dobbs, 38. Husband, feed store owner. Chest pain. Elwood Schmidt, doctor.*
> "This whole side of my chest hurts, Elwood. It really hurts."
> "What about your heart--any irregular beat?"
> "I haven't noticed any. Elwood, I just want to feel good again."
> "That's a reasonable request. And I think it's very possible you will."
> "But what do you think? Is it my heart? Is it my lungs?"

"Now, you won't believe this--but I don't know. I do not know. But I wonder. Are you lifting any sacks down at the store?"

"I lift some. But only fifty pounds or so. And only for the woman customers."

"I think you'd better let your lady customers lift their own sacks. If I know those ladies, they can do it just as well as you can. Maybe better."

There are other problems which are derivative from the "consultant" model. Perhaps the most obvious one is common sense. The argument runs as follows. When we talk about building a medical expert, a legal expert, an expert in any domain, we already are formulating the problem in an improper way. You can't be a medical expert without also being a person. There is a substratum upon which medical expertise is built and it involves in some way, growing up, going to school and living in society, knowing about feed sacks, etc. To talk as though you can compartmentalize a human being, and carve out one piece, and capture that piece in a computer program is naive, aside from the fact that it is demeaning to humans. You can't build expert systems until you have computer programs that are like children, and then like adolescents, and then like young adults, and then like medical students. Without this development, the foundation for expertise, the human "world view," will be lacking.

There are areas, however, in which we have been able to factor expertise in the sense suggested above. The inventory control person I mentioned earlier has been automated out of a job. It was not important to know whether he enjoys picnics, likes to read, likes a good joke or likes a beer. We have been successful in restricting the domain of the application sufficiently to allow the automation of the expertise in question. It remains to be seen, however, whether we can extrapolate from this experience and say that a general practitioner or a lawyer is just this problem written a little bit larger. Perhaps we have to circumscribe a problem area that may be a factor of ten or one hundred larger, but still will be able to ignore whether our doctor likes a beer once in a while.

This is the issue of real world knowledge, not just real world knowledge about the domain of expertise, but real world knowledge about people, society, and the world as well. At present we know of no "thought" experiment we can perform to assess the scope of the problem. We need to attempt the construction of some experts even to understand the problems clearly.

Knowing What You Know

As long as I am talking about knowledge, let me mention another problem we have to face, the problem I am going to call with some reluctance "self knowledge." That is, an expert knows what he knows. An expert is able to say "This is my knowledge in this domain." Now there's some question, in fact, as to whether things are as straightforward as I've just indicated. It may be that an expert only has vague understanding of what he knows, but he can roughly describe his knowledge to you (i). If you go into an automobile repairman and tell him that your car doesn't work right, he can tell you in general what he knows about cars like yours. That's important because in the interaction between consumer of expertise and the expert, much depends on explanation, the explication of the relevant knowledge.

First, consider problem acquisition. In order for the expert to convince you that he understands your problem, he must recount your problem to you. He may cast it in a slightly different light, but he's got to recount your problem so you are assured that he's understood you. Therefore it is important for him to be able to explain the knowledge that he has and perhaps the knowledge that he's going to apply.

There's also a pragmatic reason for this need for self-knowledge and explanation. Expert systems are not going to spring full blown and perfectly operational from some laboratory. In fact,

----

(i) The emphasis here is important. Few experts can state what they know with sufficient clarity and precision to allow the listener to understand it as the expert does.

most of the expert systems we can imagine are going to be capable of making rather gross mistakes. The problem here is that the user must know if the system makes a mistake. Now if the mistake is sufficiently bad, it's not much of a problem. For example, let's say we build a program that is a kidney specialist. It believes that you have a germ affecting your kidney, so it suggests that your kidney be put in boiling water for an hour, because that will certainly kill the germ. Obviously we would say, "You stupid system. That's idiotic!" It is very clear that the system is wrong. Let me give you a different example, however, which is more to the point.

A program was written at MIT some years ago by Charniak (i) that was an expert system in a small domain, the solution of calculus word problems that have to do with rate changes. These problems, on the surface, represented a fairly rich variety, and the program was impressive, because of the diversity of the problems that it solved.

Consider the problem: A barge whose deck is 10 ft below the level of a dock is being drawn in by means of a cable attached to the deck and passing through a ring on the dock. When the barge is 24 ft from and approaching the dock at 3-4 ft - sec, how fast is the cable being pulled in?

The program's answer was wrong! Now we have an odd situation. We have eighteen problems in a row correctly solved by the program, but this one is solved incorrectly. It's the right solution to the wrong problem. The boat, according to the program, was pulled up the cable! The program had no concept of gravity. If you think a simple patch will correct this, I will then give the program a problem about a kite being pulled along on a string. Charniak drew from this a lesson on the importance of real world knowledge. That is, you can't solve these problems unless you know about boats and kites and gravity.

- - - - - - - - - -

(i) Charniak, Eugene, CARPS: A Program Which Solves Calculus Word Problems, MIT Masters thesis, July 1968.

My message in using this example is slightly different. If your program doesn't have a good explanatory capability, you will never know that such a mistake happened. If you are going to check the system "by hand," or rely on intuition to check the answer, why have the system? Why not do the problem and the system can check you? If a system makes occasional mistakes, you had better hope that they are really gross mistakes in the absence of an explanatory capability of some sophistication.

Compounding this explanation problem is the necessity of finding terms that the user can understand. A program for a certain application may solve the problem, but if it can't explain to the user how it solved the problem in appropriate terms, the user may reject the solution. There are procedures that have been certified, and we accept their results without explanation. One needn't explain how the calculator computed the logarithm of a given number. You accept the answer on faith, because the calculator has been certified. For important areas of application, however, no system will achieve such reliability, and no comparable certification will be possible. Hence, the need for explanation will be great.

Optimism or Pessimism?

Even from this cursory discussion, it should be clear that there are formidable problems confronting researchers in the field of expert systems, problems that constitute major impediments to the mechanization of computer-based consultants. We face these problems with some enthusiasm, however, because we are armed with a growing "technology" for explicating, representing, and disseminating certain kinds of knowledge via computer systems.

This capability is relatively new, dating from the late 1960's, and arises mostly from research in artificial intelligence (i). In addition to technical advances in programming, there have been substantial advances in theories of learning, language, and common sense reasoning. On the other hand, a certain degree of skepticism might be voiced concerning the wisdom of such enthusiasm, because over the past twenty years various groups have claimed that such capabilities were either already accomplished, just "around the corner," or probably impossible!

The test of whether these advances will provide sufficient leverage for a fruitful attack the problem of expert systems will be in the construction of of examples. Only from such a test will we learn which, if any, of these ideas are important and lasting. Rather than consider these new ideas in detail, I will describe what I believe is an important ramification of their acceptance and use by computer scientists.

I noted above that our theories of expertise are rudimentary at best and that one of the first problems that we must face is the development of more refined notions about the basis for expertise. In this context, I commented on the inadequacy of unaided introspection in the task of explicating expert knowledge. Further, the experimental methods of psychology are not sufficiently powerful to deal with such complexity. Cognitive psychologists have made relatively little progress toward understanding expert behavior, because there was a serious shortage of ways to describe the more *procedural* aspects of that behavior. Here I believe that the advances in computer science alluded to above can play a fundamental role.

- - - - - - - - - -

(i) See for example, "New Progress in Artificial Intelligence," Patrick H. Winston, (ed.), Artificial Intelligence Memo AI-TR-310, 1974

As has been argued (i):

> The community of ideas in the area of computer science makes a real change in the range of available concepts. Before this, we had too feeble a family of concepts to support effective theories of intelligence, learning, and development. Neither the finite state and stimulus response catalogues of the behaviorists, the hydrolic and economic analogies of the Freudians, or the holistic insights of the Gestaltists supplied enough technical ingredients to develop such an intricate subject. It needs a substrate of debugged theories and solutions to related but simpler problems. Computer science brought with it a flood of such ideas, well defined and experimentally implemented, for thinking about thinking; only a fraction of them have distinguishable representations in traditional psychology.

With these tools in hand, we can contemplate new approaches to the study of expertise. Theories of expert problem solving, framed in computer science terms, can be subjected to meaningful experimentation. In this way, the computer becomes a key in the development of computerized expertise. By this I mean that this development requires a good theory of expertise, but that theory is most apt to be developed through the use of the computer as an experimental tool.

The increasing availability of such tools coupled with the rapidly decreasing cost of computer hardware allows the consideration of what I would call "unaesthetic" theories of cognition. By this I mean that in the past there was a tendency for researchers to seek parsimonious and aesthetically pleasing theories for such abilities as learning, language and the like. Theories which were "messy" in their reliance on large collections of facts or many small procedures usually were rejected either because of the the lack of aesthetic appeal or because of the lack of the requisite technical capacity for implementing and testing these theories. With new programming languages which facilitate the representation of a variety of procedures and types of knowledge, and with rapidly growing computer memories, new, more complex and multi-faceted theories of cognition can be entertained.

(i) Marvin Minsky and Seymour Papert, Artificial Intelligence Progress Report, January 1, 1972, Artificial Intelligence Laboratory, MIT.

Expert Systems and Artificial Intelligence

I want to consider the difference between artificial intelligence research and research on expert systems. After all, expertise is problem solving, and with explanation, natural language and the rest, aren't we really in the domain of artificial intelligence?

It is my contention that artificial intelligence is what I would call "theory driven." Researchers there select an application, a problem domain, that they feel will give them the most leverage toward the development of a particular facet of the current theory of intelligence. No particular credit is given for an area which is socially useful. In contrast, one can undertake work that is "application driven." We come upon some problem area, and decide that we've got to do something to improve matters, regardless of the techniques or theories we find useful; it is constantly the application that matters. This latter posture is the one assumed by expert system researchers.

Being driven by the application in complex problem domains, however, we are led into the domain of artificial intelligence, but with a rather pragmatic view. We want to solve the medical problem, the legal problem, the automatic programming problem, but in order to do that, we need a better theory. In medicine, the more that we investigate the problem of building a medical expert, the more we confront the major issues that are current in artificial intelligence research, but we see them from a different point of view. We eschew investigations that do not seem to relate directly to our application, whereas, artificial intelligence researchers would work in medicine only if it were thought to contribute to the particular theory that was in question or the particular issue that was being debated at that time.

Some people in artificial intelligence used to believe that the deep study of any particular application was a mistake, because it led just to "programming" (as opposed to leading to insights into intelligence). Burrowing like a mole into some area, digging and digging, was not considered to be aesthetically pleasing or intellectually satisfying. But this view is dying out. Thinking about thinking in the abstract has proved misleading at best. One has to have an example, and if much is to be learned, the example must be somewhat "meaty."

In most of the domains studied in artificial intelligence, the person building the program is also an expert in the problem domain. In areas such as vision, language, programming, and the like, artificial intelligence researchers are, for all intent and purpose, experts. Thus, a large part of their analysis of the domain in question arises from introspection. Let me contrast that situation with the one we face in medicine. There is a major cultural gap between the people who are trying to develop the system and the people who are the experts. This leads to a different emphasis in our approach. We place a great emphasis on protocol analysis, because we can't rely on introspection (i). This is true in the automatic programming area as well. An expert in systems design is not an expert in consulting or the management of large organizations. The problem solving processes of the latter must be studied in depth before a useful system can be built.

Thus, although there have many concerns and problems in common, artificial intelligence and research on expert systems should be seen as two different, but complementary fields.

Conclusion

At the outset, I identified an inadequacy in the supply of expertise of various forms available

----

(i) Many information systems failed because business data processing people sat in their offices and said "What would I want if I were the vice president of this company?" Introspection from that level about what it's like to be a vice president often red in the construction of worthless systems. The same would happen if I said "How would I behave if I were a doctor?" My education in being a doctor is watching Marcus Welby!

to society, and tried to show that current means for expanding the supply of expertise can have only a limited impact on this problem. The search for more radical, more innovative solutions to these problems leads us to consider the computer as a repository of replicable expertise which can be disseminated to the various points of need within society. As I indicated, however, there are significant problems which must be solved before such a grand vision can be achieved.

Despite the complexity of the problems facing us, it seems to me appropriate that we now begin to work on the development of expert systems. By confronting the problems of knowledge acquisition, representation and dissemination within the context of specific problem domains, we can expect to learn much, even if in the short run we are denied our ultimate goal. The study of the expert, (in his natural habitat as it were) has already given us new insights into the nature and application of knowledge. By the same token, this study has given us glimpses of a rather large and uncharted territory, the exploration of which is a task of the highest intellectual interest which will draw extensively upon the concepts and theories of artificial intelligence. Further, there will be a rich feedback from this work into artificial intelligence research.

Despite the undeniable fascination of some of the pieces of the puzzle of expert knowledge which we are currently studying, we must admit that the overall picture remains as yet unclear. We cannot say with any certainty whether within the near future such expert systems can be built. Certainly much more work on the problems discussed above will be required before such systems can be constructed. The enormous potential value of computer-based expert systems measured by society's need, however, argues for an immediate attack on these problems. The work which has been conducted has already yielded certain important insights. With time, continued effort, and new resources, even more progress is to be expected. It is my belief that the rapidly advancing state of computer technology and rapidly growing social needs require that we make this investment.