AI Magazine Volume 4 Number 1 (1983) (© AAAI)

The Current State Of AI: One Man's Opinion

Roger C. Schank

Yale University

AI IS IN A SERIOUS STATE OF DISRUPTION The field has always been somewhat fragmented as a result of the history of its development. Different laboratories have tended to keep their own best students, and have inbred to the point where the theories and programs invented at one laboratory have often had very little relationship to those developed at another If this parallel development had involved only theories and programs, the problem might not have been so acute. However, it has usually been the case that even the idea of what topics are significant to work on has developed in parallel. AI would have a fairly difficult time justifying itself as a scientific discipline with a coherent developmental pattern Without a coherent methodology and clearly defined goals, sometimes it seems that all AI people really agree on is that computer models that replicate intelligent behavior are worth working on (in some cases).

Recently, this haphazard situation has gotten worse. It is natural that many AI people have a need to create programs that are used by somebody for some real purpose This need has begun to express itself in a way that may alter the nature of the field profoundly. The advent of the business world's interest in AI, together with the frustration that most AI people feel about legitimizing what they do in the eyes of the world have changed the face of AI, but, unfortunately, this change has not necessarily been for the better.

In this article I wish to address some of the problems that confront AI. I am giving, no doubt, what amounts to no more than one man's opinion. It is my hope, in expressing these opinions, that the issues begin to be discussed in some public forum. I will attempt to start this debate by answering some questions about the field that have been posed to me over time. In some cases, what follows are questions that I have simply posed to myself.

General Issues

What is AI all about? In general, I see two possible answers to this question. First, AI can be seen as a modern methodological tool now being used in the ancient enterprise of the study of mind. This is the **scientific answer.** Although I do not believe that AI is a science yet, it is the intent of many of its researchers to create that science. Thus, the first answer is that AI is concerned with finding out how people think, and that, by the nature of our discipline, we demand a process-based response that can, in principle, be made to work on a computer. That is, we are looking for explanations of human mental processes that either involve, or lead to, algorithms that replicate those processes.

The second response, the **technological answer,** is that AI is an attempt to create a certain new technology That technology is concerned with getting computers to do a range of things that are quite fantastic given the current technology. However, only a small fraction of advances in computer software and hardware qualify as advances in AI. The tech-

nology that AI people want to create usually involves solving some fundamental problem the nature of what kinds of things are part of a computer program. It also usually means getting a machine to do what previously only humans have done before (rather than simply improving existing techniques). The problem with this definition has been obvious to AI people for some time As soon as something radically new has been accomplished, then since computers have done it, it is no longer uniquely human, and thus no longer AI So, one question that needs to be answered on the technological side is, can some definition as to the nature of AI software be made such that, under all circumstances, it will be seen as uniquely part of or derivative from, AI?

Why do AI? There are really only three reasons to "do" AI From the scientific point of view, you should do AI because you are interested in the mind From the technological point of view, you should do AI because you want to create new things that will be of use to society. The only alternative to these points of view is to do AI because it is fun. That point of view is what allows people to work on computer chess or computer music. (There are, of course, exceptions to this. One can study chess or music in an effort to learn about the mind, but it is not always the case.) Of course AI can be fun, but notice how few people are working on the above two subjects. Judging from the people I have encountered in AI, this is not because there is no interest in those subjects.

What are the issues? Researchers in AI have a problem They must decide what to work on. Now, to some extent, this problem exists in any academic enterprise But in AI it is a more complex problem. Most workers in other fields know what will constitute an advance. In AI we are uncertain of where we are, and hence of where we want to go A large part of good work in AI has been just answering the question of what the issues are. How many parsers or edge detectors or theorem provers do we need until we realize that these programs must be used in concert with other programs, in any system that performs a real cognitive task? It will not do to create each piece of a hypothetical system in serial fashion. Each piece has an effect on other pieces Until we identify, at least roughly, what the entire system looks like, it is difficult to make great progress on a subpart. We may find that the need for the subpart has been obviated by the apparatus that must operate on top of it.

Establishing what the issues are is a serious problem in AI A tremendous number of papers have been rejected from the recent AI conferences as **junk.** The program committees of the last three years have bemoaned the poor state of AI research. But no one quite knows what to propose as a solution to the problem since people have not been able for the most part to identify what the problem is. In my opinion, the problem is that AI is, for a set of unrelated reasons, a field without serious leadership at the moment. What is needed is a clear identification of what the important issues are, a discussion of these issues, and a clear research plan that identifies, as an aid to new workers in the field, what constitutes an advance.

For my part, I would propose that the following are the issues of the day:

1. **Integrated systems.** How do entirely integrated systems work? In both language and vision there has been a tendency to build small parts of systems. It has become increasingly difficult to know how to judge the value of such subsystems. It is important that AI concentrate on building complex integrated systems that make some stab at doing a total walk through of some real task, such as face recognition or language translation

2 **Knowledge acquisition.** There has been a great tendency lately to exploit particular technologies by adding yet another piece of knowledge to them It reminds me of my days in linguistics as a graduate student In those days it was possible to get a Ph.D in Linguistics for applying the latest technology, in this case Transformational Grammar, to yet another language. Large numbers of foreigners could be seen in any American linguistics department doing a Ph D. by writing a Transformational Grammar of their native (and hopefully obscure) language.
The same phenomenon seems to be occurring in AI these days The analogy can be direct, as in a parser for French or German, or indirect, as in an expert system for production of paint. In neither case do we have scientific excellence. It isn't even technology of a serious sort It is applied AI and should be treated as such

3 **Dynamic modification.** AI went through a long period of trying to find out how to represent knowledge We needed to find out what was learned before we could even consider working on learning itself But, most of us have always wanted to work on learning. Learning is, after all, the quintessential AI issue What makes people interesting, what makes them intelligent, is that they learn People change with experience The trouble with almost all the programs that we have written at Yale is that they are not modified by their experiences No matter how sophisticated a story understander may seem, it loses all credibility as an intelligent system when it reads the same story three times in a row and it fails to get mad, bored, or even to notice Programs must change as a result of their experiences or else they will not do anything very interesting To be blunt, an AI program that doesn't learn (and by learn here I mean become different as a result of its own actions) is no AI program at all.
In fact, I will now give a definition of AI that will disqualify most of its practitioners. AI is the science of endowing programs with the ability to change themselves for the better as a result of their own experiences. The technology of AI is derived from the science of AI and is, at least for now, unlikely to be very intelligent. But, it should be the aim of every current AI *researcher* to endow his programs with that kind of *dynamic* intelligence.

4 **Theoretical basis.** There will always be a desire

of many people in AI to create formal theories that provide some system to what would otherwise seem like sets of random heuristics. The dispute between these formalists, and more intuitive researchers, has been referred to by me (elsewhere) as the **neat/scruffy distinction.** This distinction is not all that important really. It is more or less an internal dispute in AI. A good theoretical basis is, of course, a worthwhile pursuit in AI.

5. **Creativity.** Scientists and technologists would both agree that what is most fascinating of all is the possibility that computers will someday surpass human beings They are most likely to do this by being creative in some way Principles of creativity, combined with the other powers of the computer, are likely to create this ultimate fantasy To this end, I believe it to be necessary for AI people to become familiar with work in other fields that bears upon this issue Other issues such as consciousness and development relate here also Thus, another issue is relating ideas in AI to those in allied fields with the purpose of coming to some new scientific conclusions

6. **Relationship to reality.** One general issue that seems to be coming in AI is the problem of relating what we build to the reality of what we are trying to model, for example brains, eyes, ears and so on Linking our research to work in other fields on such issues can be quite significant For our part, we may learn some things about the nature of what we are designing On the part of others in different fields, they often do not understand the functional roles of the entities that they are dealing with. Talking with people who have thought about these issues from an AI perspective might help

7 **Tools.** As always, new hardware and software tools are necessary These constitute part of AI if they are designed, not to just solve the needs of AI researchers, but to solve some of the issues of AI

Thus the scientific answer to what issues there are is the sum of research in the above areas The technological answer to what the issues are, is, of course, different From the point of view of technology, we must establish what needs to be created Building a system that no one would ever use, and there have been many in AI, serves no purpose At this point we are learning little from such systems Unless someone wants it therefore, and unless it really works, it is not a piece of AI technology. It is a piece of AI technology if it draws upon prior scientific work in AI and meets the above criteria

Who cares about what AI people create? This is a more important question than most AI people realize For years now, AI has owed its existence to giant ARPA grants that have allowed us university researchers to pursue what we pleased for the most part In fact as large companies have recently gone into AI, researchers whom they recruited were told that they could continue to do whatever they liked, and these AI people believed them. We are these naive ivory-tower residents But the people who are funding AI expect results At ARPA, ONR, AFOSR, NSF and so on, there are many enlightened people who really just want to support good science. But, even in those cases, the supporters of AI must defend their support to others higher up in the bureaucracy. Often this means Congress or some applications-oriented Admiral.

Why do I mention all this? Because it is only a matter of time before AI gets asked to put up or shut up What should we put up then? Here, it depends upon which answer (above) you subscribe to. The scientific answer says that what AI people must put up are valid, potentially verifiable theories of the mind. These theories should, in the best case, bear upon some other enterprise, such as psychology, philosophy, neurophysiology, linguistics or some such field. The problem with this answer is that, for the most part, the above fields are not very heavily supported by either government or industry. If AI were really just an attempt to do more science, then AI most certainly would have been supported in the manner of the above fields, that is, hardly at all. We have been supported in much more lavish style because everyone—theorists, government sponsors, industrial entrepreneurs and so on—has a definite belief in the essential reality of the technological answer. It may not be bad if some professors, some naive researchers and a lot of students believe that ultimately we are not all in trouble if we don't produce some technology that someone can actually use. However, the field as a whole had better not believe this. We all depend on the sponsors of AI being satisfied with its progress.

Until now, we have not had to face this reality, but all the recent publicity, investment of capital and such have forced the issue. Thus, the answer to "who cares?" must relate to the technological answer. That is, it is the people who plan on using the technology who care. AI has incurred a debt to produce some new toys that the outside world can use This may include the military, business, or just average Joes But one way or another, all of a sudden everyone cares what we do. Our field is in peril if we forget this new facet of our existence

Technical Problems

What are the important technical problems that must be solved? No matter how right or wrong the above worries/prognostications turn out to be, there are still some very hard problems facing the field that must be solved before any serious successes can be claimed. Some of my favorites are:

1 **Explaining Failures** - Any serious AI program should be able to make predictions about how events in its domain will turn out When these predictions fail, which they certainly must in any realistic system, an intelligent program should not only recover from the failure, but it must explain the failure. That is, programs must understand their own workings well enough to know what an error looks like and be able to correct the rule that caused that error in addition to being able to recognize that situation when

it occurs again. As an example of the kind of thing I am talking about, a computer should be able, by use of the same basic scientific theory, to do an adequate job of forecasting stocks or weather, or playing a game of chess or coaching a football team. What I mean by **the same basic theory** is that the theories of prediction, recovery from error, error explanation, and new theory creation should be identical in principle, regardless of domain.

2. **Representation Theory**. Probably the most significant issue in the standard AI problem of the representation of knowledge is the problem of automatically altering those representations. The problem is that these structures, no matter how adequately formulated initially, must change over time. Understanding how they are changed by actual use during the course of processing information is one of the major problems in representation itself. Deciding when to create a new structure or abandon an old one is a formidable problem. Thus, new AI programs should be called upon to assimilate information and change the nature of the program in the course of that assimilation. Clearly such programs are necessary before the knowledge acquisition problem can be adequately attacked It should also be clear that an AI program that cannot build itself up gradually, without requiring all its knowledge stuffed in at the beginning, is not really intelligent. Thus, here I am thinking of a program that developed opinions and ideas greater than those initially given it, such as a foreign policy specialist, or an editorial writer that functioned by reading the news regularly and altering its opinions as the facts warrant.
3. **Search** I believe, as do many other AI people, that search is one of the key AI problems. However, I think that the approaches to search have been inadequate. Searching massive amounts of information requires not efficient algorithms but representations that obviate the need for those algorithms For example, to refer to a practical problem, consider the problem of building a travel agent who knows everything there is to know about hotels in resorts. Finding what you need for the right customer is not only a question of search, it is first necessary to establish exactly what is being searched for. This is a far harder problem than the search itself. The issue is really how to represent the knowledge or, more precisely, what knowledge to represent.
4 **Reconstructive memory**. Psychologists have known about the phenomenon of reconstructive memory for some time. People are capable of answering questions from very incomplete data. They can figure out if they should know something and whether they might be able to figure it out. Such self-awareness depends strongly upon an ability to know how the world works in general—or, the representation problem again. Building a program that would know if it would know is a very important task
5 **Generalization** A program that can form a testable generalization from experience would be of great significance This program would have to be able to draw conclusions from disparate data The key aspect of a good generalization maker is his ability to connect together experiences that are not obviously connectable. This is the essence of creativity Thus, we might consider building models of political leaders, or criminals, or stockbrokers, attempting to automatically formulate theories based upon the data available about what these people are likely to do. These theories would be based upon an ability to generalize from prior events. These generalizations would likely be inadequate at first, but eventually new theories that fit the data should emerge.
6. **Reasoning from cases**. Ultimately human expertise is embodied not in rules but in cases. People can abstract rules about what they do of course, but the essence of their expertise, that part which is used in the most complex cases, is derived from particular and rather singular cases that stand out in their minds. Thus, I am proposing expert systems that are basically sets of unusual cases The job of the expert is to find the most relevant case to reason from in any given instance Thus, an automated judge might be a good first attempt here because of the case-based nature of the law However, all reasoning programs that give advice (something I am fond of calling Advisory Systems) should eventually behave in this way if they are to handle the hard parts of expertise.

Which of these are most important? All of them are important of course. But one thing above all: an AI program that does not learn is no AI program. Now, I understand that this maxim would not have made much sense in the past But one of the problems of defining AI, is as I have said, that AI could, by past definitions, be nearly anything. We have reached a new stage. We have a much better idea of what is learned, therefore it is time to demand learning of our programs. AI programs have always been a promise for the future, a claim about what we could build someday. Each thesis has been the prototype of what we might build if only we would. Well, from the technological perspective, the time to build is now. From the scientific perspective, after the issue of what is learned is taken care of, the issue for AI is learning, although we probably do not have to wait for the former to be finished in order to start.

What constitutes a Ph.D. thesis? A Ph.D. thesis is a promise for the future. Nobody should get a Ph D for building yet another piece of technology. There should always be the real belief that the thesis has yielded another piece in the puzzle. In other words, an AI thesis should, at this stage in AI, tell us something about learning. It should not tell us an even faster algorithm for parsing or another area where rule-based systems can be applied.

The Real World

Where does AI fit in Computer Science? AI has happily been a part of Computer Science now for twenty years or so. But, it should be pointed out, Computer Science

has not always been all that happy with AI AI played a role in the inception of some of the major Computer Science departments, but other departments have seen fit not to have any AI or to have just a token effort. Why is this the case?

AI has a rather poor reputation in Computer Science. I will not argue here whether this is justified by the evidence. The major question is exactly where AI belongs. Part of the problem here relates to the issues I discussed above with respect to what constitutes a reasonable piece of AI research. Consider for example, a fanciful but not impossible program that generated folk tales from a universal base The algorithm would be the same regardless of the culture and language that were input to it as data. The program would then generate folk tales relevant to that culture. Assuming for the moment that such a program were possible, would this be a valid piece of AI research?

I would like to argue that as interesting as it would be in principle, most AI people would not take much interest in it. This would be the case for a variety of reasons, not the least of which is the fact that most AI people know or care little about issues that are essentially anthropological My point is not so much that this work would not be AI (it would be) but that research is only deemed significant if it is appreciated by some scientific group. AI people wouldn't appreciate it, but anthropologists would It would, in fact, be a very significant piece of anthropology

Why do I mention this? There will come a time, I predict, where this problem will become the rule rather than the exception We are already beginning to see some of this in the expert systems work. AI should, in principle, be a contribution to a great many fields of study. AI has already contributed some to psychology, linguistics, and philosophy as well as other fields Really what AI is, potentially, is the algorithmic study of processes in every field of inquiry. As such, the future should produce AI/Anthropologists, AI/Doctors, AI/Political Scientists and so on. There might also be some AI/Computer Scientists, but on the whole, I believe, AI has less to say, in principle, to Computer Science than to any other discipline The reason that this has not been so heretofore is an accident of birth. AI people have been Computer Scientists, therefore they have tended to contribute to Computer Science. Computer Science has needed tools, as has AI, and on occasion these tools have coincided AI is actually a methodology applicable to many fields It is just a matter of time until AI becomes part of other fields and that the issue of what constitutes a contribution to AI will be reduced to the question of what constitutes a contribution in the allied field At that time what will remain of AI will be precisely the issues that transcend these allied fields, whatever they may turn out to be In fact that may be the best available working definition of what constitutes a successful contribution in AI today, namely a program whose inner workings apply to similar problems in areas completely different from the one that was tackled originally.

What prevents AI from getting out into the world? The U S. has made a great many mistakes with respect to capitalizing on its scientific advantage and converting it into a long-term technological advantage. Innovators in computer science are up for grabs in the marketplace, professors are leaving for industry, or for foreign shores. The problem is one of the nature of the interrelationship between industry, government and the university, but in a sense it starts in the university.

The modern university encourages entrepreneurship in its professors. A professor with innovative ideas in the sciences must have laboratories and assistants In another era, the university could provide for the relatively smaller demands of scientists, but this is no longer possible. Keeping the books balanced preoccupies university presidents far more than attracting great scientists. The university administration may want scientific superstars but it doesn't want to have to pay for them Enter the professor who can pay his own way.

Once upon a time, professors were members of the upper class. One way a professor could pay his own way was out of his own pocket. Today's professor is more likely to be from the middle class. It is the United States Government or a private foundation who pays. Now the question is, how does a professor get them to pay and what are the consequences of a professor getting a million dollars a year to spend?

The answer to these questions can be summed up in one word—entrepreneurship To survive in the modern university, a scientist must become an entrepreneur in the fullest sense of the word. He must learn where sources of money are; how to convince the controllers of that money to give it to him; how to manage the people and facilities that that money will buy him; and how to get his fair share from the University. The only thing the entrepreneurial professor does not do is make a profit.

But what happens when the research of this entrepreneurial professor begins to bear fruit? The University is no place for applied work, so the government is placed in the awkward position of having to stop supporting the work just as it nears fruition. The reason for this is that there are many disparate sources of government money Those that deal with universities are not the same kind as deal with private industry Often a whole new set of contacts has to be made, and a long hiatus can occur. Further, the professor has no place to do his applied work (that is, if he is interested in doing it at all).

But to some extent, the health of the overall enterprise requires that the applied work be done People expect to see new technology when they have been supporting it for a long time with their hard-earned money

Where is private industry in all this? To some extent they are always willing to grab up new ideas but, in AI at least, private industry has been a reluctant partner. Lately a few companies have become interested in AI, but there may not be enough practitioners of AI to staff those companies. To be able to do AI, and to be willing to do applied rather than basic research, requires a type of person whose numbers are less than plentiful.

One question here is why there are so few AI people. Why hasn't AI been able to expand very easily? The answer is complex. By and large most Computer Science departments have been hostile to AI. They tend to believe that AI has not been able to do what it claims, which is, for the most part, true. Also, other computer scientists tend to feel threatened by AI. But the major reason is that, in general, AI people have not wanted to go alone into a new environment without the right machinery and without colleagues who understand their enterprise. AI is very difficult to do, even more so if you have to do it alone.

How does business relate? Business relates to AI in a great many ways. It will encourage the training of applied AI researchers, by employing them when they are produced. It will provide a way for AI research to get out into the real world, thus keeping sponsors and critics appeased. It will provide some direction for what AI people should be worrying about, but therein lies the rub.

While AI should worry about producing things people might want, that being what AI is about from the technological perspective, it cannot allow itself to become absorbed in those issues to the exclusion of the serious and unsolved technical and scientific issues I believe that that is unfortunately what is happening today There is, for example, a difference between technological completion and technological innovation.

A lot of what passes for AI research lately has had very little in it that was clearly innovative. For example, we are seeing new expert systems in this or that without attempts at solving the theoretical problems that relate to expert systems These include having your expert system learn from experience by changing its rules when they fail and having the expert system naturally "grow" and reorganize his initial knowledge over time. Clearly these issues are the hard scientific problems that underlie expert systems, and while they need not be worked on by everybody, they are what constitute research I have no objection, indeed I applaud, the attempt to bring expert systems to the market place if they are ready The two questions that I feel compelled to ask are: (1) are they in fact ready? and, (2) do AI people believe that completing a technology by finding new ways to apply it, and changing a technology by doing research, are the same thing? The first question I address below. As for the second question, let me assert that scientific research and AI research ought, in principle, to have no conflicts. Sometimes what goes under the rubric of AI research is not research at all, but application of research.

In some sense business has been affecting AI for some time. In AI we have always been aware of the eventual utility of our work, and, despite work on computer chess and computer music, we have also seen lots of AI work with obvious applications in the military Business only differs from the military in its particular application, not in its overall effect on research AI without practical import is, in a sense, what we have been calling Cognitive Science. That is, when we do work simply for the scientific knowledge that we gain, we are merely using our particular methodology to solve ancient problems concerning the mind.

What business will do, in the long run, is focus some of the issues that we work on in AI toward applications that relate to the real world as opposed to those that don't For example, one can address the search problem in a variety of domains, some practical and some not. If the effect of business on AI is to force theoretical work within domains of some practical relevance, then I believe that to be a good effect. Orienting projects toward real problems will actually force us to deal with realistic issues rather than fanciful ones. AI found out a while back that chess was not the only or best source of interesting problems. I see no problem with realism unless it disallows work on problems that have no obvious practical realization.

One question to be concerned with here is the extent to which the success or failure of AI business entities will affect the field as a whole. AI is in the prosperous state that we see now because there are many jobs available and much research money available. If these two things dry up, AI will be in dire straits. The question is whether or not failure of the current AI ventures will cause the currently advantageous situation to change. I believe that it would, drastically I do not believe that ARPA would be able to survive the pressure from the outside world if business were to believe AI to be a failure. I also think that universities would cease to want AI people In other words, I think the entire future of AI as a field depends upon the success of the AI companies. Of course, we can always muddle along like philosophy and English literature, in the event of massive failure, but as we are newer, our situation is likely to be worse, reduced to an effete subject remembered for its quaintness.

On the other hand, if AI companies succeed, the field will have other problems. All of a sudden there will be a need for AI masters degrees and for more PhDs. AI will be something people study to get jobs (to some extent this is already the case). We will become established and stodgie. In either case, AI as we knew it in the sixties and seventies, a bastion for purists worrying about essentially insoluble problems, may well be gone for a long while.

Does AI matter? In some sense, all subjects of inquiry are really AI. All fields discuss what the nature of man is; AI tries to do something about it. From a technological point of view AI matters to the extent that its technology matters, which is still debatable But from a scientific point of view, we are trying to answer the only questions that really do matter. Personally, I think technology is great, but let's not have it at the expense of the scientific enterprise.