AI Magazine Volume 6 Number 3 (1985) (© AAAI)

The Dark Ages of AI: A Panel Discussion at AAAI–84

Drew McDermott
Yale University, New Haven, Connecticut 06520

M. Mitchell Waldrop
Science Magazine
1515 Massachusetts Avenue NW, Washington, D C. 20005

Roger Schank
Yale University, New Haven, Connecticut 06520

B. Chandrasekaran
Computer and Informational Science Department, Ohio State University, Columbus, Ohio 43210

John McDermott
Department of Computer Science, Carnegie-Mellon University, Pittsburgh, Pennsylvania 15213

Drew McDermott:

In spite of all the commercial hustle and bustle around AI these days, there's a mood that I'm sure many of you are familiar with of deep unease among AI researchers who have been around more than the last four years or so. This unease is due to the worry that perhaps expectations about AI are too high, and that this will eventually result in disaster.

To sketch a worst case scenario, suppose that five years from now the strategic computing initiative collapses miserably as autonomous vehicles fail to roll. The fifth generation turns out not to go anywhere, and the Japanese government immediately gets out of computing. Every startup company fails. Texas Instruments and Schlumberger and all other companies lose interest. And there's a big backlash so that you can't get money for anything connected with AI. Everybody hurriedly changes the names of their research projects to something else. This condition, called the "AI Winter" by some, prompted someone to ask me if "nuclear winter" were the situation where funding is cut off for nuclear weapons. So that's the worst case scenario.

I don't think this scenario is very likely to happen, nor even a milder version of it. But there is nervousness, and I think it is important that we take steps to make sure the "AI Winter" doesn't happen—by disciplining ourselves and educating the public.

This panel has been assembled to discuss these issues. I've asked the panelists to discuss the following questions in particular: Are expectations too high among consumers of AI, such as business and military? If they are too high, then why? Is there something we can do to change this mismatch between expectation and reality? To what extent is this mismatch our fault? There's a charge often leveled against AI people that they claim too much. To what extent is it due to näiveté on the part of the public? What is the role of the press in this mismatch, and how can we help to make the press a better channel of communication with the public? What is the role of funding agencies in the future going to be as far as keeping a realistic attitude toward AI? Can we expect DARPA and ICOT to be stabilizing forces, or is there a danger that they may cause people in government and business to get a little bit too excited? Are funding agencies going to continue to fund pure research, even if AI becomes a commercial success? Will the perception remain that we need to do some things that are not of immediate commercial interest? And, finally, what should each of us do to insure his survival in case of problems?

Here to discuss these issues are Mitch Waldrop, from *Science Magazine*, representing the press; Ron Ohlander, from DARPA, representing a funding agency; Roger Schank, from Yale University; B. Chandrasekaran, from Ohio State; and John McDermott, from Carnegie-Mellon University. The first speaker will be Mitch Waldrop.

Mitch Waldrop:

First, I would like to relate an experience I had earlier this week when I was attending a seminar in New York state that Isaac Asimov organizes every year. This year the topic was Artificial Intelligence, and the idea was to bring in people from all walks of life and, over the course of several days, work up a human impact statement for artificial intelligence. Marvin Minsky, as well as myself and others, were on the resource panel. The result was what you might expect: A combination of silliness and seriousness, with not a great deal of informed insight into AI. But there was a very good cross-section of the general public, and I gained some very interesting insights while trying to answer their questions.

One, is that most of these people make essentially no distinction between computers, broadly defined, and artificial intelligence—probably for very good reason. As far as they're concerned, there is no difference; they're just worried about the impact of very capable, smart computers.

Enthusiasm and exaggerated expectations were very much in evidence. The computer seems to be a mythic emblem for a bright, high-tech future that is going to make our lives so much easier. But it was interesting to hear the subjects that people were interested in. Education seemed to capture their imagination most—computer-aided instruction potential. If you want to see some real passion, start talking about what happens to people's kids in their school-room environment.

This was followed by an absolute fascination with cognitive science and what artificial intelligence is telling us about how we think. As for applications to health—they were a little vague beyond potential for diagnosis. In fact, they didn't make much distinction between artificial intelligence and biotechnology.

There was even some interest in the possibilities of what could be done with very large databases, searching it with intelligent database searchers, etc.

I'm not sure what it means, but it's interesting that this seems to be in roughly the inverse priority to what AI people give these subjects.

What really struck me was the flip side of exaggerated expectations—exaggerated fears. The computer is not only a mythic emblem for this bright, high-technology future, it's a mythic symbol for much of the anxiety that people have about their own society. The most obvious, what you might call the "1984 Big Brother Is Watching Anxiety," is that somehow the computer will erode our freedom and invade our privacy. Who writes the computer-aided instruction programs for our children? They control what our children learn, and how they think. There seems to be an implicit assumption that there's always going to be some big, manipulative power structure up there controlling things.

A second anxiety, what you might call the "Frankenstein Anxiety," is the fear of being replaced, of becoming superfluous, of being out of a job, and out on the street.

A third, closely related anxiety might be called the "Modern Times Anxiety." People becoming somehow, because of computers, just a cog in the vast, faceless machine; the strong sense of helplessness, that we really have no control over our lives, that computers, being inevitably very rigid, brittle machines, becoming more and more powerful, inevitably result in alienation, isolation, enforced conformity, standardization, and all those bad things—leaching away of humanity. I'm going to come back to this, but I'll leave it right now by saying that these fears and expectations are not groundless.

That brings me to the news and the media. My colleagues in the news and media have heard me rant about imbecile reporters and relentlessly shallow TV reporters, and I'm not going to give that speech here. It would be superfluous.

But I bring up the general public's attitude to point out that reporters, editors, and TV people, are human beings; they are reflections of the society in which they live. They write about or film what their readers are interested in and what they are interested in; what seem to them to be important issues. That brings us to the key problem in covering something like AI. The problem is not a matter of imminent deadlines or lack of space or lack of time or people straining for "gee whiz" or "Oh, my God" type headlines. The real problem is that what reporters see as real issues in the world are very different from what the AI community sees as real issues, and the trick is to bring these into consonance. Where does that leave the AI community? There seems to be broad agreement here that the coverage of AI is abysmal. So what do you do about it? Something that is very unhelpful is to take an attitude that everybody out there is a pack of idiots except us, who really understand. Cheap shots at reporters get big laughs at meetings like this but are not very helpful. Some of us HAVE taken predicate calculus and can understand it very nicely.

What would be useful is to ask yourselves: If you don't like the coverage that you're getting now, what would good coverage consist of? What would be a good story about AI? What would you like to see? AI is about giving operational definitions for things. What's an operational definition of a good story? I'm not going to attempt to answer, but it might be helpful if you think about that.

When a reporter comes to talk to you, what message are you trying to communicate? It might be helpful to know ahead of time. If what you want to communicate is the latest stuff on some nonmonotonic, backward chaining, I don't think it is going to be too compelling to the reporter. If it is about the nifty expert system that you hope to be marketing next month, it's probably going to look very funny to the reporter. He would be very suspicious. But just in general, ask yourself what it is you are trying to communicate. Think of it as an opportunity, not as an interruption and an irritant.

The idiot reporters and the insensitive editors are always going to be worthless. I cannot offer you any wisdom or magic formula to make them go away. The best you can hope to do is help the conscientious reporters who do come by and may just be confused or not know, but are genuinely willing to try to learn as much as they can given their time constraints. These people do exist. And as I said, it helps to have a clear idea of what you want to accomplish.

A final point is a modest suggestion to the community to get out in front of this Frankenstein issue. We're going to hear a lot on this panel about overheated expecta-

tions. I think the fears are just as strong, and they're out there in the general public. There's no point is dismissing them as neurotic or misled. They're there, no matter what the source. There's no point in saying it is not AI's problem, that it's robots who are kicking people out of the factories. Well, you guys are designing vision systems for robots. People are going to be thrown out of work by various forms of expert advisors, at least temporarily, and people are not going to like it when their careers have been dashed by some superanimated Visicalc. Governments can use large databases to violate people's privacy and to harass them. For that matter, credit card companies can do that, too. One can even envision a natural language system that monitors telephone conversations. Computers and even AI programs can be made to be rigid and constraining. There's no point in saying that you don't need computers to do all this. It is true that Hitler managed to create a totalitarian society with no help from computers. But computers can also be used to exacerbate these tendencies and a power drive for human stupidity. If we expect physicists to be concerned about arms control and chemists to be concerned about toxic waste, it's probably reasonable to expect AI people to be concerned about the human impact of these technologies.

Just to make sure I am not misunderstood: I'm not advocating that people run around spouting leftist rhetoric or crying "chicken little." I am suggesting some sober and constructive thought about how one can order, how one can address these problems over the long run, and that the community as a whole take some kind of position. After all, computers, especially with the aid of artificial intelligence, can be made extremely flexible, extremely personalized, and can produce a great deal of wealth—we hope.

I'll conclude by saying that perhaps if people had something substantive to say on these issues, reporters would not have to strain for "Gee whiz" or "Oh, my God" headlines.

Roger Schank:

I'm sorry, I'm not representing the business interests today. I hate to disappoint you. But in line with that, I'll tell you that I have a new company that does educational software. I mention it because I was having a conversation with Oliver Seldridge, telling him about my new educational software company. I said, "Well, it really doesn't have anything to do with AI at all, except that some of the software we design has to do with things like teaching reading and reasoning comes from ideas that we've had in AI, but there's no AI in the programs in any way." And he said, "Oh, sort of like expert systems, huh?"

I came here to relay to you six conversations. That's the first. They're all short.

The second conversation I had was with a real estate developer, who had a Ph.D in Biology. He wanted to build an expert systems industrial park; every company in it would be doing expert systems. I said: "You may have come to the wrong person; I don't much believe in expert systems." "How can you say that?" he said. I asked, "What do you mean?" He explained, "Well, to get computers to model everything that somebody knows; to put all the knowledge in and have the thing be just like a person—that's terrific." I replied, "Yes. But we don't know how to do that yet." He said, "No, but that's what expert systems are." That's the second conversation.

The third conversation was with Bob Wilensky, a former student of mine, who asked me what I was going to do on this panel. He asked if I thought doomsday was coming. I said, "Yes." And he said, "No, you're wrong." I asked why. He said, "It's already here. There's no content in this conference." Now I think there's something serious to be concerned about there. He isn't the only person I've heard express that view. If that's true and there's no content in this conference, then doomsday *is* already here.

Conversation four was with a representative of ARPA—not Mr. Olander—and he said, "You know, we've got a real problem. We've got so much money to spend on scientific research that we have more money than there are scientific researchers." And I replied, "And we're not going to be able to fix that, are we?" It's very hard to make more scientific researchers in an environment where money is being offered in tremendous amounts to be developers rather than to be researchers. ARPA can't raise salaries. It can only offer money to hire people.

The fifth conversation isn't a conversation. It's just a report. It's a report of what I heard was the standard for accepting papers to this conference this year. I've been on the program committee a few times, but missed it this year. What I heard was that only completed scientific work was going to be accepted. This is a horrible concept—no new unformed ideas, no incremental work building on previous work. I don't know if that's actually what happened. I didn't attend a lot of sessions. I can tell you that if that is what happened, that's frightening. Remember AI? See, you guys may not be as old as I am AI-wise. But I remember AI, the first conference and the second conference. We used to sit and argue about things—not whether or not we should go public. There were always people in AI who were interested in software development tricks. That's great. There has always been that component in AI and there always should be. But if it comes to dominate AI, then AI becomes applied systems. I don't like that.

The last conversation was with Eugene Charniak, and I had it ten years ago. He kept saying to me, "Roger, you're promising too much. You can't do all the things you think you want to do. They're very hard." And I said, "Yes, but they're fun, and I want to work on them. And, anyway, I think I can do them in ten years." Here's why I mention this. It's been ten years. I haven't done them. Gene Charniak is a wise man. I may not be able to do

the things I thought I could do ten years ago in fifty years. Yet, at the same time, as I'm beginning to discover more and more problems about why things are hard, we are getting less and less people working on those hard problems and more and more people working on applied situations. That's frightening.

What do I think the issues are here? First, I think the press is completely and utterly irrelevant. When I first got into this field twenty years ago, I used to explain to people what I did, and they would say, "You mean computers can't do that already?" They'll always believe that. And it doesn't matter what the press believes, and it doesn't matter what the general public reads in *Time* and *Newsweek*. It really doesn't make any difference. We have responsible reporters in *Science* magazine, but not that many people read *Science* magazine in the general public. I don't think it's an issue what the general public believes. However, it is a very important issue what big business believes. You see, big business has a very serious role in this country. Among other things, they get to determine what's "in" and what's "out" in the government.

I got scared and started a company at the same time, when there were lots of startup companies around. I got scared when big business started getting into this— Schlumberger, Xerox, Hewlett-Packard, Texas Instruments, GTE, Amico, Exxon,—they were all making investments— they all have AI group. You start to wonder who could be in the AI groups. We haven't got that many people in AI. And you find out that, those people weren't trained in AI. They read an AI book, in many of these cases. They started off reading all the best AI research. After a while you discover AI group after AI group whose people were only peripherally in AI in the first place. What's going to happen is that those companies will find that their groups aren't producing as well as they had expected. When they find that, they will complain; they will say nasty things about AI. The presidents of those companies will be talking to the people who are not at ARPA, but at the Secretary of Defense level. They'll say things like, "Well, I've spent so many millions of dollars on AI this year, and I've had it. They're not producing anything." And it may be that AI is capable of producing things. It may be that even the people at those companies are good. But it also may be that it'll take them more years than anyone expected.

I'm very concerned about this issue. It's the reason I'm on this panel. And I think it's very important for people not to worry about the press. Talk to the press; they're nice. It doesn't hurt.

The thing to worry about is when you hear that a company is starting up an AI effort, you better ask: When do you expect what? The small companies, the startup companies, that have been started by AI people, don't present the same problem. We've had to learn how to build a product. The more we learn about products, the more we begin to realize that our products look less and less like AI. That's okay—that's what derivatives are like. It's okay to build derivative things. What's not okay is to build *only* derivative things. That's frightening. So I am concerned not that expectations are too high, but that expectations are too low. What expectations am I talking about? The expectations that we, as an AI community—I assume that the people left here at this conference are the actual AI community—have forgotten that we are here to do science, and that we are nowhere near the solution. We used to sit and fight about these things in public; now we all sit and talk about how it's all solved, and we give slick talks with slides with pretty pictures on them. I'm very concerned about the fact that people don't want to do science anymore—it is the least appealing job on the market right now.

It's easier to go into a startup company and build products. It's easier to go into a big company and have a little respite and do some contract work. It's easier to do all those things than to go into a university and try and organize an AI lab, which is just as hard to do now as it ever was, and sit there on your own trying to do science. It's difficult. But if we don't do that, we will find that we are in the "dark ages" of AI.

I take this opportunity to talk to any of you who are considering the choice between the two and recommend that you seriously consider that the science of AI is critically important and also, by the way, a lot of fun.

I leave you with two messages, which will be obvious, I hope. The first, from one half of my life, is incumbent upon AI because we have promised so much, to produce. We must produce working systems. Some of you must devote yourselves to doing that. And part of me devotes myself to doing that. It is also the case that some of you had better commit to doing science. Part of me commits to doing that. And if it turns out that our AI conference isn't the place to discuss science, then we better start finding a place where we can discuss science, because this show for all the venture capitalists is very nice. And I hope all the people back there sell more computers and more systems, and they should all live and be well. But I am concerned that people here who are first entering this field will begin to believe that a Ph.D means building another expert system. They're wrong.

Ron Ohlander:

They should always put Roger Schank on last because nobody can top him.

I'd like to go back to the possible scenarios that Drew McDermott outlined. I think there are some other possible scenarios. I'll mention them, and what the government is doing concerning them.

One possible scenario is that the current upsurge of development continues unabated, and we have most remarkable development going on for the next few years. Another possible one is that in the near future we're going

to have some shakeout and realignment. A third one is we'll have some disillusionment with the process. There will be some retrenchment. One way of characterizing it is we would possibly go back to the state of affairs that were in existence four or five years ago. Finally there is the last scenario that McDermott described. And, of course, there are all the things in between.

I will leave those scenarios for now and discuss some of the things that are happening in government and how they portend for the possible downside of the AI technology.

The government has increased its interest in this particular area. DARPA, has been involved for quite a number of years in AI research. But the other services are also starting research efforts. DARPA is continuing its basic AI program, and it's still funded at the same level with some sliding piece over the next few years. In addition, there's the strategic computing program, which was described to you this morning. The Army has started a support program for an AI laboratory to be placed at a univeristy. The Navy has established an AI lab at NRL. The Air Force has started to put together an AI consortium at the universities to support them in research and development and in education and training. In addition, the intelligence community is taking a long hard look at AI and what it offers to their endeavors.

What I conclude from this is that there's a strong, healthy interest in AI, that there's not a lot of wild-eyed people out there that expect things that are beyond the state of the art, and that most people are taking a very orderly approach to the exploitation of AI within the government. The people that are heading these efforts are pretty level-headed—they know what is going on. I think that government funding is likely to continue. It's hard to get started, but once it gets started, it's likely to continue for a period of time.

In addition to these efforts, there are a number of things going on in terms of system development. This gets into the actual exploitation of AI for particular systems. In case anyone thinks that the situation of AI is unique, let me tell you I have worked for three or so years at Naval Electronic Systems Command, and I looked at practically every development that went through there involving computer technology, and every one of them had trouble. The fact was there were a lot of failures. There were overruns and systems delivered past schedule. This is certainly not unique to Naval Electronic System Command. The most would be systems being acquired for the government. The government continued to acquire such systems simply because there were also a lot of successes—key successes. This kind of success encouraged continued interest and development in the field.

The fact is, there aren't enough good people to go around, so the government and everybody else are forced to pay high prices for successful systems. These high prices not only come for particular developments, but they come for multiple efforts to get the same development because of the failures that occur. I think the government interest will continue, and that we will see the systems as development contractors do more work and get more interested; they will make representations for the government for incorporation of a lot of the technology into these systems. The implementations will occur and some successes will also occur. There will also be some failures, but I think we'll live through it.

I'd like to make one other comment. What kind of steps can be taken to circumvent the dark ages or to decrease the impact of the downside of the current interest in AI? I'm not sure that a lot can be done. We're pretty much riding the whirlwind.

There are a few conclusions I would like to draw from my observations of what's happening in the field; my observations in funding various efforts. Let me repeat: I think there will be a lot of failure. On the other hand, I think there will be enough key successes to override the failures to keep interest focussed on the area. The shakeout will come in the not too distant future. There will be some shakeout in machine areas and some shakeout in all the expert system technology companies. I think the government will sustain their funding. In fact, it's likely to increase as people get more interested in applications. And in some rebuttal to Schank, I think there's room for both scientists and engineers in this field. What we're seeing is the rapid growth of an industry that has no underlying engineering support, so scientists, who would otherwise be doing research, are filling that role. It is also my observation, as it's been Schank's, that people currently getting involved in industry are coming from in-house assets, people who are being trained. I see nothing wrong with that because I think there is a role for people to take this scientific technology and to implement it from an engineering standpoint.

So revisiting those scenarios that I postulated earlier, I think that basically the interest will continue, and we'll have some shakeup and realignment. In summary, my outlook for the future is rather positive.

B. Chandrasekaran:

As one of the few academics who is not in modern business, I have been assigned the role of survivalist. First I want to ask, "Has AI paid its way?" ... Or to put it another way, "Have we earned our keep?" I have three answers to that: Yes, yes, and yes.

It's been the most profound paradigmatic change, the most performed paradigmatic change in human understanding of some of the important issues about ourselves in a long long time.

The notion of cognition as computation is going to have extraordinary importance to the philosophy and psychology of the next generation. And for well or ill, this notion has affected some of the deepest aspects of our self-

image. I think it's for well, but we'll have to see. Even at the technology level, what we have been able to do is minimal if you measure it against the capability of the human mind. But if you measure it against changes in the styles of programming or styles of building things, significantly useful accomplishments have been made. The problem is not that AI is weak in terms of its usefulness or importance. AI has been the whole expectation of the problem.

When people start building complicated things, there is a remarkable consensus of what they are looking for. The reason for that is already there is a software architecture based on frames; the embedded procedures and moving around in that kind of space constitute a weak theory of mind, but nevertheless a theory of mental architecture. The next is very minimal commitment, which is for the good, because we don't know much more at this point to make strong commitments. In that sense it is already helpful and will enable us to build entities that we wouldn't have been able to build otherwise.

The next question to ask is "Has there been a discontinuity in AI to justify the sudden interest from whatever viewpoint?" Yes, there has been discontinuity, and the discontinuity has been in the idea that knowledge is very important. We can't separate knowledge as something that is pragmatic and go away and do theory that is not connected with knowledge. So in that sense, the discontinuity that has caused this interest is the importance of knowledge. It's true in natural language understanding. Roger Schank and his group's work and several others have emphasized that. Problem solving has also been important. The problem has been, however, that we have very weak theories of knowledge and even weaker theories of how to use them. People don't understand that very well. So they're being mistaken and misunderstood with respect to their power. Most importantly, there have been no characterizations of what is possible. So the real problem is the very, very strong belief in omnipotence of simple architectures.

There is also a real confusion because the computer science community, the AI community, has refugees from so many areas. Symbol-level theories, which may even be right, are being mistaken for knowledge-level theories. This is one of the conceptual problems that has been bedeviling us.

So, basically, the problem isn't underestimation of the problems of multiplicity of generic knowledge structures. Knowledge use invalidities has been the cause of misunderstanding. We're seeing some systems and extrapolating that all it takes is more of the same. I also believe that faster architecture could do the trick. Let's run 30,000 rules. Let's run 50,000 frames. The idea being that all it takes is more architecture, faster systems.

We need to characterize the things that we can do—that's do-able. Then spin off a list of things that we know how to do and let people go ahead and do them. We could also more clearly understand the research issues that need study.

Regarding the commercial prospects; that may no longer be a very interesting problem from an AI viewpoint, but it would not have been possible without AI having been there. This class of problems can be categorized as knowledge-rich and problem solving-poor. All that is needed is getting some form of knowledge, organizing it, and making it available to people. It's going to have very little problem solving capability, but without the recent history of AI, such a thing would not have been possible. But it's not going to solve all the problems. A reasonable number of problems can be handled with this appoach. The formula I normally use, 10% AI and 90% other, is what is going to make it useful. In fact, this might even spin off, and in some ways it may even be better off for AI if it does. Then AI can get back to concentrating on research issues.

The problem in applying even current technology is that they still require epistemic analysis. This is very hard. To think that all it takes is engineering is a mistaken notion.

Epistemic analysis is hard to teach. Some people can use the same tools and build extremely interesting systems; other people cannot. So AI is going to be blamed for the failures of people who do not have the capabilities of sufficiently powerful epistemic analysis. Also, AI is going to be blamed for what I call dilution of AI. We started from AI, which then became expert systems, which then became rules, which then became LISP. So people think they are doing AI when they do LISP programs, for example. It is important to keep reemphasizing that AI is not all those things. AI is something else. We must keep emphasizing the importance of all those "something elses."

Will there be an AI Winter? I think there may be an AI dusk, which may be even better, rather than the hard sunshine that we have been having in Austin recently. Instead of the bubble bursting, the bubble may become somewhat smaller and less fragile. So it may be an interesting place to be. I think AI has already contributed enough, and I believe that AI will contribute enough to justify itself at that level.

There are all sorts of historical analogies about what's going on in AI. One relates to the automatic high-quality natural language translation. I don't think that analogy is valid. I think that it was based on too small a number of ideas. It didn't have enough robustness and solidity to it. AI today has more robustness and solidity to it in so many levels that that's not the real analogy. A truer analogy is probably closer to biotech, where even five years ago people thought incredible claims were being made.

But biotech companies have not gone bust. Many them have gone back to solid research. Even five years ago, AI people used to go around talking about how they're going to clone the human mind. That used to get me really

scared because that showed that big theories were being mistaken for strong theories.

My hope is that AI will evolve more like biotech in the sense that certain technologies wil get spun off, and researchers will remain and extremely interesting progress will be made.

With respect to projects such as the Japanese fifth generation, there is a nightmare of some bureaucrat at the DARPA finally taking all those things and stamping "Failure," "Failure," "Failure," on each one of them. Our position ought to be to plan for success and to realize that in many ways we cannot fail. If we have some personal and professional integrity in what we do, we cannot fail to come out ahead. Even at a technological level, enough things will be happening for DARPA to get its money back. That's the sort of attitude to encourage. That requires not pushing weak theories too far.

What am I doing as a surviror? I believe the next four or five years are going to be some of the most exciting years in research. In spite of all the over-promising and over-expectations, the last five years have identified extremely important and interesting problems. I hope to be involved in that kind of research. Also, unlike some other areas, we don't have to decouple technology from research.... You can build things and still do research. In that respect, AI is in a better situation than other so-called theoretical sciences. As long as we characterize each of the advances in knowledge-level terms, identify what it's capable of doing, and identify what it's not capable of doing so that users know what kind of problems can be solved, I think we can come out in reasonably good shape at the end of this period.

John McDermott:

I want to revisit Drew McDermott's original thinking on why this panel is a good idea. His position is that there are these things called expectations. The people who have false expectations are going to become angry, upset, or unhappy with the violation of their expectations, and walk away. That applies to our situation, if the people of the world find out that AI isn't what it expected, then AI isn't going to be supported anymore. Presumably, one of the kinds of support that we need most is funding for our research efforts. So that's going to go away.

Schank offered a slight variant of that, which is that we don't have to worry about what the people think, but we do have to worry about what big business thinks. If big business gets angry at us, then the funding source will be cut off, and we won't be able to do science.

I think that we all agree that we would like to have the science of AI continue to be supported. I don't think anybody, either up here or out there, is at all unclear on the fact that we haven't yet made much progress in the field, and that we've got a long way to go, and most of the exciting discoveries are still ahead of us. So we somehow want to insure that the funding base for the science of AI continues to be there so we can do good research.

What bothers me about Drew McDermott's premise is that it's not completely clear to me that people who have violated expectations are going to end up withdrawing support. If we look at the kinds of expectations that could be incorrect, there are presumably many. I have four that I would like to go through. If you focus on different types of expectations and the ways those could be violated, it's hard for me to see a clear connection between violating those expectations and not having research get funded.

The first kind of expectation that could be off would be an expectation about the kinds of tasks that current AI systems or AI systems five years from now, would perform. I have encountered people who have a science fiction view of the world and think that computers now can do just about anything. But that view never seems to manifest itself in a personal way. These people have a feeling that computers can do wonderful things, but if you ask them how exactly could an AI program help in work, they don't have the sense that within a week or two they could be replaced or that computers can come in and do a much better job than they do in their work. So I don't think that at a concrete level, people have a näive view that super-intelligent computer programs are right around the corner. In fact, I find the opposite. When people talk with me about systems that could be developed to solve particular problems, I'm often the one that describes a more positive or grandiose role for the systems than they do. The people who might fund these efforts end up having what seems to be an extremely healthy caution in what they expect. So I don't see a lot of wildness in people's ideas of what systems can do. Even if there are some people who do have overblown expectations along those lines, I think that the AI technology has developed to an extent now that it is possible to produce software that can do some extremely helpful things. Because of that, people are going to be happy. They might not be as happy as they would be if their wildest dreams were satisfied, but the slope is upward, and I think they'll like it.

So if someone's expectations are violated because he or she gets some good, helpful thing, but it's not as good and helpful as hoped for, I don't think that's the kind of expectation violation that's going to result in the funding rug being pulled out from under us.

The second kind of expectation has to do with the background people have to have in order to be able to produce or to build AI systems. One thing that we all have a tendency to say about AI systems is that they are easier to build than more traditional programs. Often we end up forgetting to say "easier" and say "easy" to build. It is conceivable, and I suspect that some people have translated "easy to build" into "with a one-day tutorial in AI, I can go out and build an AI system." I think that's mostly

wrong.

Let's say that a lot of businesses have some middle managers who are convinced that all they need in order to build an AI lab is to hire five or ten people, send them each out for a one-day tutorial, give them each a book, and at the end of that one day, they'll go off and'll start producing a great system. I'm hard pressed to believe that management is somehow going to get so caught up in that myth that they are going to be able to then react violently when they find out that those people don't produce as much or as quickly as they had hoped. That kind of judgment doesn't seem to me to be the kind of judgment about how long one might have to spend in order to acquire aptitude to do some task. The notion that all you need is a day or week of training isn't the kind of judgment that people make about most tasks. There might be some wishful thinking, and people might try to build a group because they can't hire trained AI people. I would expects that management has a notion of the risk and that they understand that they are taking a path less likely to lead to success, at least quickly, than a path that was built on a stronger experience base.

The third kind of expectation is that every attempt to build a knowledge base system will succeed gloriously. If the builders of the system are people with absolutely no experience in AI, I don't have any sense at all as to what might be a reasonable success rate to bet on, but I suspect it's low. Even if the people had significant experience in AI, we still know that there are many attempts to build AI systems that go on for a while, and then the people discover that they can't build the system that they thought they could. I don't think that's the kind of expectation violation that's going to pull away support for AI research, because as long as there's some reasonable amount of success, as long as some fraction of the systems that are attempted turn out to be truly helpful, there's a positive, forward-moving attitude. And people say, "Gosh, I bet I'll be more fortunate next time."

The fourth kind of expectation is the expectation about what the level of performance of a successful AI system is going to be. There may be people who believe that the successful AI system will never make a mistake. We all know that the most we could expect from an expert system is a performance level that's as high as the experts, and experts, of course, make mistakes. Presumably, we would be delighted with a performance level that was substantially less than that. So there are going to be systems being used that make mistakes. If the people who are using the systems or asked for the systems to be built started out with the expectation that those systems would never make mistakes and then find, much to their surprise, that they do make mistakes, that's an expectation violation. But again I don't see that that implies, that that suggests in any way, that these people are somehow going to turn away from AI. They are going to become better informed. They are going to understand that the technology that they are dealing with has certain limits, and they will presumably be in a position to understand why those limits are necessary.

If you ask yourself what are the kinds of expectations that people might have and how might those expectations be violated, it seems to me that the nature of the violation is going to be such that the people who had the false expectation will simply become better informed about the nature of AI technology. The nature of the violation is not going to result in rejecting the technology, feeling cheated, or feeling that somehow AI doesn't have any promise, with a result of these people turning away and putting their research dollars into something else.

Questions from the Audience

Audience:

Following up on Dr. Schank's suggestion that big business is the driver in this expectations failure model or the AI Winter, I was wondering if people on the panel could identify any operationally successful expert systems that have been implemented over the last five years. I know John McDermott can tell us about R-1, but that is fast becoming ancient history. Where are some systems, as John suggested, being implemented that are being viewed as being successful? I don't see any.

John McDermott:

Let me try to restate the question. There is, at most, one successful AI system out there, and if AI has any promise, why is that? What are the names of some of the other successful AI systems?

When I give talks on successful AI systems, I put up slides that have systems with names. Some of the slides that I put up refer to systems like ACE, CATS1, the Drilling Advisor—there are a number. If somebody says, "Just how successful are those systems? To what extent are they being used on a regular basis? How much money are they saving? And so on..." I'm not directly involved with any of those systems. I don't know how the companies that participated in the development of those systems are using them and whatever. I'm told that the systems are getting better and that the people who work with them think they are valuable and see promise in them. I do have occasion to go to Digital Equipment Corp. from time to time, and in addition to the R-1 system, Digital has four or five other systems that, though they have been used for a shorter period of time and aren't used as widely. Nevertheless they are being used. There's a system called XCEL, which a large number of sales people are now beginning to use. There's a system called IMAX, a fairly simple system, a part of which is now being used in a manufacturing plant.

There are AI systems out there, and the people who use them say, "Gosh, I'm glad these systems are here."

And I suspect that this phenomenon is going to continue. I believe that the technology is able to provide a lot of assistance, but it's taking us some time to develop those systems. The transfer of AI systems into real working environments is going to take a few more years. In a few more years when somebody is asked to name systems, the list will be much longer, and it will no longer be an issue of trying to find enough systems so that we're not embarrassed.

Audience:

As has been pointed out several times, it's clearly true that the explosion in applications work has siphoned off a lot of our precious Ph.D level talent for AI research. It's also true, I think, that it has created an entirely new resource, and that is what you might refer to as journeyman level or master's level AI people. Not all those who have very little or no formal AI background are not ex-Ph.Ds. Not all these people are of the sort who read one book, go to a one-day tutorial, or spent three days working on AI. In fact, in some of the industrial labs, you have people who do have, say, a Master's in Computer Science, who started out perhaps by reading one book and going to a tutorial, but over a year or two or three of work in the field, have actually developed some capability in the field. Now these are people who never could be and never would be researchers. They're not going to get a Ph.D. But they are capable of being the equivalent of our lab technicians. The question is: Do you see any way to use this new resource in our Ph.D style research to alleviate some of the lack of Ph.Ds?

Roger Schank:

I think that it's wonderful that those people are being created, if they are. I'm a little anxious about how well they are being created, but I think it's important. So don't misunderstand me on this statement. They're not researchers. The worry I have is that they will begin to think they are researchers. I don't think we should make that demand of them. My concern about AI, by the way, and about the complexity of it and why it'll take so long, is not so much that there is a tremendous lacking of ideas that we need, but also that there is so much engineering involved. We have tremendous amounts of knowledge to put into these systems or else get them to learn them on their own. Either way, there is a tremendous amount of engineering involved. So we do need to have this engineering class built in. I don't think AI research is very easy though, and I think some Ph.Ds aren't very good at it. So I'm concerned about the Masters people doing it.

Audience:

I'm addressing this question to Ron Ohlander. The question deals specifically with the history of speech understanding research and the funding of DARPA for that, but I think it reflects the general problems with funding for AI. From 1970 to 1975 there was a project aimed toward continuous speech understanding undertaken by DARPA, and it funded several centers. It was perceived as having failed. There was a severe retrenchment in the latter 1970s for speech understanding research, and now I see that under the strategic computing projects, continuous speech understanding once again is being funded with, I gather, fairly high expectations of success. My question is, what has been learned from the previous round of funding speech understanding research and also of AI, so that this time things will work out differently than they did fifteen years ago?

Ron Ohlander:

I assume you are addressing the question from a political standpoint rather than a technical standpoint. And that is, are we going to have that kind of experience again? I wasn't there at the time, so I can't tell you from my own personal knowledge, but I've talked to people. I think the problem was the casuality of the political process and that there were new things that people wanted to do. I think that program was sacrificed to some of the new things they wanted. I will also point out that the bulk of the AI program stayed intact. I think the chances of that same kind of thing happening again are nil. There is a lot more consciousness of the world of AI, and the programs are supported very strongly. In fact, we wouldn't have gotten the strategic computing program started without the enthusiastic endorsement of Dr. Cooper, who is head of the agency.

Audience:

AI appears to be a fairly young field. I'm sure that most people in the room over thirty or so are fully aware of this, but as a word of kindness for those who are young and starting in the area, most of our high technology areas go through boom and bust cycles. Think back through aerospace industry, back when Sputnik kicked things off. I'm a mathematician. We had a cycle in the early 1960s when we could not produce enough Ph.Ds. Many kids started training, and then in the late 1960s the market was glutted. Many Ph.D mathematicians couldn't get jobs in mathematics at all, much less in a university at the Ph.D level. This happens in all fields. The universities try to control it as best they can. So do the industries. But these things happen. You also appear to be on a slight boom period with a potential bust coming. And there will be survival. That will happen. I agree with that. The field will exist. It will grow. It will be robust. But there will be many of you in this room or your students who may get shaken out in the long run. As a touch of kindness to you, there are several things that can be considered. Check with the other disciplines. We have been through this before. In the MAA Journals and AMS Journals for Mathematics

in the mid-1960s, around the 1965-1967 time-frame, have very serious articles discussing this and how to get out of it. IEEE transactions for engineers did the same thing. So did the aerospace industry. Someone jokingly said: Find what else to call yourself. For those in AI, a very easy bailout will be as a computer scientist. If that fails, for those of you that are appropriately trained, about five years from now, or seven, or ten, if I read the statistics correctly, call yourselves mathematicians. We'll need you again.

Audience:

I'd like to ask Roger Schank a question concerning the flavor of this convention. I'm a student, and I came to this convention with certain hopes and expectations, and the flavor is a little bit different than I expected, to say the least. What would we do to change the flavor of IJCAI-85 or AAAI-86? What could we do to change this to gear it not so much toward business and more toward the research? Or more toward getting back to arguing and less about selling?

Roger Schank:

Well, you're not going to change this. AI conferences have an evolution, and if you've been to all of them, you begin to see the process. I can see that this one is going to get more business, not less. The issue is to be able to start other forums in other places where people can start to have those kinds of dialogues. We've tried to do that from time to time. It requires a certain amount of energy for somebody who actually wants to run a conference. But I think that that's probably what's needed, that AAAI ought to concern itself with running more than one conference, and it ought to have one where there aren't any booths and there aren't any tutorials, and see how many people show up. Maybe the people who showed up would want to talk about science.

Audience:

I hate to pick on Professor Schank, but I have another question directed at him. If corporate centers can't really produce researchers, isn't it the responsibility of an organization like AAAI to provide more tutorials during the year where people from corporate research centers can and get more information and become closer to Ph.D level researchers or at least not lead corporations down the garden path, thinking that they are actually doing AI work?

Roger Schank:

There's a presupposition in your statement that is wrong, that through lots of tutorials you'll learn to be an AI researcher. I think that's blatently false. AAAI should go on the way it is. Tutorials and shows are wonderful. But if that's the only conference we have, that's the problem. We just don't have the other ones. With respect to training people, Ph.D research, at least in my laboratory, used to take three to four years. Now it seems to be taking five to six. It's a long process. You have learn a lot of stuff and then try to create something on your own. That's what a Ph.D means. All I can say is that that isn't the same thing as being trained through tutorials. There's a master's degree program at Stanford, but I don't know how many other AI master's programs really exist. Probably what we need are more of those, where people in industry can take off for a year and learn about AI techniques and then return to their company. I don't think tutorials will ever do it because if you don't have hands-on experience, it doesn't work in AI.

John McCarthy (from the floor):

I'd like to comment on Schank's point in the question to him. The AAAI does sponsor workshops, and anyone who wishes to organize a workshop in some special topic should get in touch with me, because I have undertaken to continue that.

Audience:

I have a question about historical perspective on this entire issue you're discussing. There have been a number of technologies that have run into dead ends, like dirigibles and external combustion engines. And there have been other ones, like television and, in fact, the telephone system itself, which took between twenty and forty years to go from being laboratory possibilities to actual commercial successes. Do you really think that AI is going to become a commercial success in the next ten to fifteen years, or will it have a longer gestation period? That's addressed to the entire panel.

Roger Schank:

There are differences between AI and derivatives of AI. AI means that you are going to make machines intelligent. Remember that? We were interested in the mind and what it meant to be intelligent and how thinking took place. Remember all those issues? That is what AI is about. Now that doesn't mean that the companies discuss those issues on a daily basis. Companies are great—they should exist. And they should do derivative work. That means providing software which is better than the software that we had before, which has all kinds of AI flavor, all kinds of AI derivatives. It just shouldn't be confused with AI. That's all.

John McDermott:

Let me add to that. As we learn some things in AI, it simply makes it possible to develop systems whose behavior is better in certain kinds of circumstances. And that's going to continue. So I don't see that there are any sharp boundaries. I don't know what it would mean to say that we're going to try to go back and have increasingly

dumber programs. We're going to understand better how to solve certain classes of tasks, and that's going to continue. People are going to like that, and so they're going to encourage these derivative companies to produce that kind of software.

Audience:

What I was trying to point at was that you can have a technology develop and then continue at a very low level. We still have steam turbines. We still have lighter-than-air aircraft. But they're not a major commercial force.

Mitch Waldrop:

Picking up on your comment about historical analogies and also picking up on Chandrasekaran's talk about work on hard problems, historically you can look at development of the steam engine back toward the end of the 18th century, beginning of the 19th century. This was done by people just trying to solve the immediate problems. In the long run, however, because they had to understand the nature of heat and work in steam engines, they discovered thermodynamics. Carnot was an engineer trying to understand the efficiency of steam engines. In the long run, forcing yourself to work on real-world hard problems can help.

Audience:

This is just a comment on what John McDermott had to say about companies not really having high expectations. A lot of the real tight purse strings, the big money for AI, is being directed just toward places that are gearing their people up from reading the book, going to a tutorial, and, believe it or not, those companies do have high expectations of the software that they are hoping to come up with. And I'd just like to say that dismissing that as *not* a big issue, as far as money drying up, an AI Winter, or an AI dusk, is not something we should forget about, if we're worried about funds for research.

Audience:

I'd like to say also that I've visited a number of labs of large companies that have very serious AI efforts going on with extremely large research programs. Those companies are succeeding in turning out good products. I think that will counterbalance the effect of the companies who do not seriously address the technology.

B. Chandrasekaran:

I would like to make one remark about the whole trend of using expert systems. At least it has done one thing well. It has concentrated a lot of analytic effort in some areas of the company's operation or in some body of knowledge. But at the end of it you might conclude, "Gee, I don't need any expert system. I understood the whole thing very well. I can summarize the whole thing in eighteen rules. Or I can just write up a simple system to do that." But that still says it has a good effect on the operations of the company. I think if those things take place, AI should also take some credit for it too, because it has emphasized looking at knowledge and analyzing it properly.

Audience:

I'd like to address this to anybody on the panel who wants to answer. What feelings do you have about the sentiment that AI is rather fragmented and is a bunch of solutions looking for problems and that there is very little methodology being developed to allow people to vigorously analyze problems and pick and choose the methodology to use in different parts of the problem? Right now it's a very arcane art.

Audience:

I don't think anybody would disagree with that. It's what you would expect in the early stages of an engineering discipline.

B. Chandrasekaran:

I would like to comment that most people have a false notion of rigor as proving mathematical theorems. I think a lot of people are pretty rigorous. They get up in the morning and tear their hair out over issues, and they don't accept the first solution to come to mind. But they don't express their solutions as mathematical theorems. So with that particular proviso, I think people are aware of it and people are doing rigorous work. But they may not recognize it as rigorous because people might have wrong notions about the definition of rigor.

Audience:

I'm probably one of the few business people still here. I wanted to say a couple of words of reassurance about business expectations. I come from a Fortune 50 company. We've got a small AI group of four, but we're not claiming to do AI. We're attempting to find things that we can pick now and apply. I think you'll see that creating a constant demand—a constant demand for perhaps what you're calling engineering. We're not seeking to be AI researchers; we never claimed to be. What we will become will be the moral equivalents of systems analysts. That kind of thinking we can learn from doing. And I think that that is a point that has not yet been said. There's a different kind of demand for a constant kind of application that you're going to see, and that will flatten your boom and bust cycle.

Audience:

Well, I concern myself with the following issues. I don't know who you are or what company you come from, but

I concern myself with the training of the people who are at that company, who you now claim are an AI group. I worry how much AI they know.

Audience:

Very little. And we'll be up front about that.

Roger Schank:

But soon they will be doing applications in your company, which if it's a Fortune 50 company, is an important company and will be claiming or knowing about the successes or failures of this AI group and its ability to do applications. I'm concerned whether those people know anything at all about AI.

Audience:

And what we'll say in response to that is we have found one more information-handling technology that is applicable, possibly in conjunction with others, that helps solve some problems. We're not fools. We will pick applications where we can have some incremental value. We're not going to try for a home run and strike out swinging.

Roger Schank:

It would make me a lot more comfortable if you didn't say we have a small AI group—if you said we have some smart programmers who are working on some hard problems.

Audience:

Then perhaps that's the term you would feel more comfortable with. We're already moving away to engage in our own expectation management. We're already moving away from terms such as "expert systems" toward "knowledge systems." We're trying to drop the term "AI." We're doing all the expectation management you're talking about.

Roger Schank:

I wish I could get a follow up on this two years from now and see how it worked out.

Audience:

We'll still be here.

Roger Schank:

I guess so.

Audience:

I just wanted to ask what you think people can do about avoiding the sensationalist journalism about artificial intelligence. You see a lot of articles now about machines that can think. People get a lot of expectations, especially those who don't have a technical background or expertise. They are starting to look for a lot to come out of artificial intelligence. I have a feeling that this boom-bust cycle, at least in the popular sense, could easily be fueled by current sensationalist journalism. I'm wondering what you think can be done to calm people down about the possibilities of AI and make them more realistic.

Mitch Waldrop:

In part, I have to say I can't help you. The *National Enquirer* is always with us. I can't make it go away. In part, it's inevitable. There's the phenomenon of the three-day wonder. This is more than three days, but it's the same phenomenon, that something new appears on the horizon, people get excited about it, get enthusiastic about it, and they gush. After awhile, it goes away. I will say that in my talks with lay people about artificial intelligence, I have exactly the same impression that John McDermott talked about in that they seem to have a vague idea that great things can happen. have sublime confidence in you people, that you can do anything. But when it gets down to the nitty-gritty, they tend to be pretty unimaginative and have pretty low expectations as to what can be done Referring back to this last weekend that I spent, I had to keep boosting people's ideas about possibilities, such as what a teaching machine might be able to do. I think it might be less of a problem than you really imagine. Yes, the sensationalism will go along for a while, and then it'll blow itself over when they find something else to be excited about. I'm surprised Michael Jackson hasn't stopped it already. But I don't know how tremendous an impact that will have on serious decision makers. Let me add something else, too, that I keep hearing. I'm again going to paraphrase John McDermott in terms that might be much blunter than he would prefer. There seems to be this unconscious assumption around here that everybody in the world except artificial intelligence researchers is kind of dumb, mindless, unreasonable, inflexible, and incapable of learning, and in short, behave not unlike computers. Most people I know out there are rather reasonable and sensible—it may not be as much of a problem as you think.

Audience:

It seems that academic AI people tend to blame everyone but themselves when it comes to problems of AI in terms of relationship to the general society. Charges of arrogance are traditional, but I'm concerned with a different one. It seems that there's a need for master's level engineering education in a more organized fashion than the zero start situation of so many AI groups in fairly large companies that are starting up efforts. It seems that Ph.D level people who can teach or master's level people with background can usually do much better things in terms of personal reward than teach It seems that there is a need for some kind of mechanism, either within the universities or possibly in doing an endrun, that creates an alternative institution for a master's level engineering education

in AI. I'd like to raise that as an issue, and ask the panel to comment on that.

Roger Schank:

The problem with the master's training is not that we don't want to do it. I think the problem is that everyone wants someone else to do it. In fact, representing most of the AI faculty at Yale sitting at this table, I don't want to do it. You don't want to do it. We're happy to sit there and say we should have master's students, but the dearth of good Ph.Ds are going to the universities to do research and none of them want to do these kinds of training programs. I hope they'll come, but I don't know who is going to start it. You have to get someone who is very dedicated to that proposition as being something important. When that happens, they occur. But it's very hard to get anyone to really care, even though you're right.

Audience:

Do you think the initiative is then more appropriate in coming from the industrial community to try and stimulate special programs in the universities? I know that Stanford has had a master's program oriented toward AI, and it was established partly because of very considerable financial support and promises over at least medium term rather than short term from industry. Notably, Bell Labs helped sponsor the masters in computer science, which was the reason why there was enough masters' programs to support masters in AI at Stanford. So is it appropriate for industry to push on the universities and come up with a lot of dollars to support these masters' programs?

Roger Schank:

That might help, but again you would have to get somebody who cared about it. There'd have to be enough in it. Suppose you took a guy who was starting out at a new university, and he was concerned about how was he going to get good computing facilities and attract graduate students. If they came to a person like that and said we'll give you extra money so you can do that, by the way, if you'll do the masters program, he might find himself at the center of a very nice situation. I think you'd have to find the right person, but it could be done.